



Co-Writing Assistant

Security, Privacy and Trust

A security and privacy whitepaper for Pearson Clinical customers and their IT, security and governance teams

Contents

OVERVIEW

- 1 Executive summary
- 2 What the Co-Writing Assistant is (and is not)
- 3 How it works at a glance

SECURITY AND PRIVACY BY DESIGN

- 4 Architecture and Amazon Bedrock
- 5 Protecting PII and PHI
- 6 Your data is not used to train models
- 7 Output quality and accuracy
- 8 Safety and guardrails
- 9 The practitioner stays in control
- 10 Operational assurance

ASSURANCE AND EVALUATION

- 11 Compliance and shared responsibility
- 12 Built on AWS
- 13 How to evaluate and verify
- 14 Appendices

How to read this document. Sections 1 to 3 describe what the Co-Writing Assistant is and how it works. Sections 4 to 10 set out the security and privacy controls. Sections 11 to 14 cover compliance, verification and supporting material.

OVERVIEW

1. Executive summary

The Co-Writing Assistant helps practitioners draft evaluation report narratives faster, while keeping every decision, and every word that reaches a client, in their hands. It is available through Pearson's Q-global and Q-interactive platforms and built on Amazon Bedrock. It is engineered so that the information you are responsible for is protected: your data is never used to train foundation models, and the practitioner reviews and approves everything.

This is a security and privacy whitepaper. It is written first for the IT, security and governance teams who assess the tool, and it gives practitioners, purchasers and special educators the same assurances in plain language. It sets out what the assistant is, how it is built, how it protects data, and how you can verify what we claim.

Our commitments

Commitment	What it means
Your data stays yours	Client information is never used to train foundation models.
Protected by design	PII and PHI are minimized, encrypted in transit and at rest, and access controlled.
The practitioner decides	Nothing is finalized or submitted without explicit practitioner review and approval.
Accuracy taken seriously	Output is screened by guardrails and always practitioner-reviewed; where Pearson has published instrument guidance, the draft is also grounded to that instrument. Quality is measured and improved.
Guardrails on every interaction	Automated safety controls screen what goes in and what comes out.

2. What the Co-Writing Assistant is (and is not)

It is a drafting aid under the practitioner's control, not an automated decision-maker. The Co-Writing Assistant generates draft narrative for evaluation reports, helping practitioners express results in clear, professional language. The practitioner remains the author: they direct the tool, review what it produces, edit freely, and decide what goes into the final report.

The Co-Writing Assistant does	The Co-Writing Assistant does not
Draft narrative text for the practitioner to review	Diagnose
Help articulate and structure findings	Score assessments or interpret results on its own
Adapt to the specific instrument in use	Submit, finalise or send anything automatically
Keep the practitioner in control of every word	Replace practitioner expertise or professional accountability

3. How it works at a glance

In plain terms, the Co-Writing Assistant works in three layers. The practitioner provides a request and the relevant context. Pearson-authored guidance tells the model how to draft a report narrative, and for assessments Pearson has published instrument-specific guidance for, that guidance also encodes the structure, terminology, scales and score interpretation of the instrument itself. Amazon Bedrock runs the underlying model inside Pearson's secure environment, returns a draft, and the practitioner reviews and edits it.

Instrument-specific guidance does not exist for every assessment. The Co-Writing Assistant supports both Pearson and non-Pearson assessments, including material the practitioner uploads, so it will draft for assessments outside that set. Where instrument-specific guidance is available the draft is more tightly constrained; where it is not, drafting rests on Pearson's general report-writing guidance and practitioner review carries more of the load. Section 7 sets out what this means for accuracy.

In sequence:

1. The practitioner makes a request with the necessary context.
2. Guardrails screen the input.
3. The model drafts the narrative, guided by Pearson's report-writing guidance and, where available, by instrument-specific guidance for the assessment in use.
4. Guardrails screen the output.
5. The practitioner reviews, edits and approves. No model training takes place.

SECURITY AND PRIVACY BY DESIGN

4. Architecture and Amazon Bedrock

The Co-Writing Assistant is built on Amazon Bedrock. Amazon Bedrock is a fully managed service that offers a choice of industry leading foundational models (FMs) along with a broad set of capabilities that you need to build generative AI applications, simplifying development with security, privacy, and responsible AI.

It runs in the Q-global US region (US East) for US customers. Pearson's application invokes Amazon Bedrock as a managed AWS service over private AWS PrivateLink endpoints, so requests originate from within Pearson's own AWS environment and prompts and outputs stay within Pearson's AWS account boundary and chosen region. Traffic does not traverse the public internet, and the model endpoint sits within the same region and security boundary as the Q-global application. The underlying models are accessed through Amazon Bedrock's managed library; they are not trained or fine-tuned on customer content.

5. Protecting PII and PHI

The most important fact for a security team is the simplest: the Co-Writing Assistant does not retain your prompts and outputs. They are used to generate the draft in real time and are then discarded; the service does not log or persist them. Everything else here builds on that. The design assumes prompts and outputs may contain client information and treats them accordingly.

Control	How it is applied	Basis
No retention	Prompts and outputs are used only for real-time generation; the service does not log or persist them.	Bespoke
Encryption in transit	All traffic protected with TLS 1.2 or 1.3 over HTTPS and private endpoints.	Inherited
Encryption at rest	Any data stored for processing and uploaded documents are encrypted with AES-256; keys managed in AWS Key Management Service (AWS KMS).	Inherited
Data residency	The Co-Writing Assistant uses a US cross-region inference profile that routes requests within the US region and remains within that geographic boundary.	Inherited
Access control	Segregated by customer account, least-privilege and role-based, with access audited; data is never shared with model providers.	Inherited

6. Your data is not used to train models

Client information never trains foundation models. We state this precisely, and we are equally clear about the one separate use that exists only with your explicit consent.

By default. Content sent to Amazon Bedrock is not stored or logged by AWS, is not used to train any AWS or third-party models, and is not shared with the model providers, who have no access to it. Pearson sends the prompt and receives the output, with no retention beyond processing.

With your consent, separately. Under a separate, opt-in Data Collection Agreement, distinct from the click-through license, a customer may authorize Pearson to share a subset of their prompts and outputs with Pearson's internal team, used solely to validate and improve the accuracy and safety of the Co-Writing Assistant (see Section 7). It is voluntary, can be withdrawn at any time, and does not train foundation models. In production this sharing is off by default unless the customer signs the agreement, and any data shared stays within Pearson's security controls.

This distinction is deliberate. Rather than claim an absolute, we set out exactly what happens to data by default and what changes only when a customer chooses to opt in.

7. Output quality and accuracy

Generative AI can produce text that reads convincingly but is wrong. For a practitioner's report that is a real risk, and we treat it as one. The Co-Writing Assistant manages it with two controls that apply to every request, whatever the assessment, and with the practitioner as the final decision-maker.

Screened by guardrails. Amazon Bedrock Guardrails filter harmful content and redact personally identifiable information on both inputs and outputs. This applies to every request, for every assessment.

Reviewed by a person. Every draft is reviewed, edited and approved by the practitioner before it is used. This is the decisive control, and it applies to every draft the assistant produces, whatever the assessment. The assistant drafts; the practitioner decides, remains the author, and the accountable professional.

No automated safeguard removes the possibility of error entirely, which is why human review is mandatory and why the Co-Writing Assistant is positioned as a drafting aid rather than an author or a decision-maker.

Grounded to the instrument, where Pearson has published instrument guidance. For the assessments Pearson has published instrument-specific guidance for, that guidance constrains the model to the structure, terminology, scales and score interpretations of the instrument in use, including score-range classifications and appropriate language. This makes the draft materially better: it ties the narrative to the actual instrument and reduces generic or off-scope statements. Coverage is expanding and is not universal, and the Co-Writing Assistant also drafts for assessments outside it, including material the practitioner uploads. Instrument guidance therefore improves the quality of a draft; it does not substitute for the review of one, and the two controls above apply either way.

8. Safety and guardrails

Amazon Bedrock Guardrails are configured to screen both inputs and outputs. They filter harmful content at high severity, detect and redact personally identifiable information in model output (for example names, addresses and phone numbers), and enforce a deny-list of eight out-of-scope or clinically inappropriate topics. Where Section 7 addresses correctness, the guardrails address safety: keeping harmful, sensitive or off-scope content out of the exchange.

Beyond the platform guardrails, custom rules and prompt constraints keep the tool within scope. It does not provide diagnoses or test interpretations beyond what is clinically appropriate, which also aligns with the FDA's definition of a non-device clinical decision support. The interface carries disclaimers reminding practitioners to review and edit all generated text, and that suggestions do not replace professional judgement.

9. The practitioner stays in control

The practitioner is the author and the decision-maker, by design. The tool accelerates a first draft; it does not replace professional judgement or sign off clinical content.

Human in the loop. Every draft is reviewed by the practitioner before it is used.

Edit, accept, override. The practitioner can change any part, accept or reject suggestions, and write their own text.

Authorship and accountability. The practitioner remains professionally responsible for the final report. The Co-Writing Assistant does not finalise or attest clinical content.

Audit trail. Each use is logged in Q-global's audit system with the user ID, timestamp, event type and related record identifiers. Sensitive content from prompts and outputs is deliberately excluded from these logs. Logs are encrypted and retained in line with Pearson's records-retention policy, consistent with HIPAA's minimum for such records.

10. Operational assurance

Beyond the controls already described, the Co-Writing Assistant inherits Q-global's operational security programme.

Incident response and breach notification. Pearson operates incident response and breach-notification procedures under its existing programme, consistent with its HIPAA obligations.

Data-subject rights. Requests to access, correct or delete personal data are handled through Pearson's established privacy processes. Inactive accounts and their associated data are removed after two years, and earlier deletion can be requested.

Availability and continuity. The service runs on resilient AWS infrastructure with encrypted backups. Regional continuity arrangements apply, with cross-region failover used only in emergencies.

ASSURANCE AND EVALUATION

11. Compliance and shared responsibility

The Co-Writing Assistant inherits Q-global's compliance posture and sits within the AWS shared responsibility model. Pearson operates as a HIPAA and HITECH Business Associate, with Business Associate Agreements in place, and the relevant controls fall within Pearson's SOC 2 Type II audit scope. The platform's security programme is aligned with NIST and ISO 27001 and is subject to regular third-party assessment and penetration testing. The service also inherits the assurance of the underlying AWS platform: Amazon Bedrock is in scope for FedRAMP Moderate, SOC, ISO and HIPAA eligibility. FERPA requirements are met for educational records, and a Privacy Impact Assessment completed in 2025 confirmed the feature added no new data-protection risk. AWS is the underlying infrastructure sub-processor.

Who is responsible for what

Layer	Responsibility	Owner
Cloud infrastructure: compute, storage, network	Physical and platform security	AWS
Application, configuration, guardrails, data handling	Secure build and operation of the Co-Writing Assistant	Pearson
Access, user management, consent, clinical use	Appropriate and authorized use	Customer

12. Built on AWS

The Co-Writing Assistant is built on AWS for enterprise-grade security, the data commitments that come with Amazon Bedrock, the breadth of AWS compliance attestations, and a managed service that is not tied to a single model provider. Building on an independently assured platform lets Pearson concentrate on the clinical quality and safety of the tool.

Pearson and AWS announced an expanded collaboration in February 2025 to bring AI-powered learning and assessment solutions to market using Amazon Bedrock. AWS documentation confirms that customer inputs and outputs are not used to improve the foundation models and are not accessible to the model providers, and that Amazon Bedrock is in scope for major compliance standards including FedRAMP Moderate, SOC, ISO and HIPAA eligibility.

13. How to evaluate and verify

Security and governance teams do not have to take these assurances on trust. On request, Pearson can provide the SOC 2 Type II report, a penetration-test summary, the Data Processing Agreement and the sub-processor list, and a walkthrough of the architecture with the product and security teams.

The questions below are the ones governance teams ask most often.

Is our client data used to train AI models?

No. Content sent to Amazon Bedrock is not used to train the underlying foundation models and is not shared with model providers.

How do you manage the risk of incorrect AI output?

Two controls apply to every request: output is screened by guardrails and, decisively, reviewed, edited as needed, and approved by the practitioner before use. Where Pearson has published instrument-specific guidance for the assessment in use, the draft is also grounded to that instrument. The assistant drafts; the practitioner decides. The system does not independently calculate scores or retrieve undisclosed external data.

Where is our data processed and stored?

US customers' data is processed and stored in the US geographic region boundary.

Is the practitioner still responsible for the report?

Yes. The practitioner reviews, edits and approves all content, and remains the author and the accountable professional.

What security documentation can you share?

On request, Pearson can provide a security and privacy overview, a Data Processing Agreement or Addendum, the sub-processor list, and relevant third-party attestations such as SOC 2.

14. Appendices

The following supporting material sits behind this summary.

Appendix A. Security controls matrix

A summary mapping of the principal controls and what they do.

Control	What it does
No model training on your data	Customer content is not used to train foundation models.
Encryption in transit	TLS 1.2 or 1.3 over HTTPS and private endpoints.
Encryption at rest	AES-256 with keys managed in AWS Key Management Service (AWS KMS).
Data residency	Region-bound processing (US).
Network isolation	VPC with AWS PrivateLink to Bedrock; no public internet egress.
Access control	Account segregation, least-privilege RBAC, access audited.
Audit logging	User, timestamp and action logged; no PHI in logs.
Guardrails	Content filtering, PII redaction and denied topics on inputs and outputs.

Control	What it does
Human review	Mandatory practitioner review and approval of all output.
Incident response	Breach detection and notification under the existing programme.
Compliance attestations	SOC 2 Type II in scope; HIPAA Business Associate.