CISO Assurance Series

Preservica AI – Security, Privacy, and Trust Overview

About the CISO Assurance Series

The Preservica CISO Assurance Series provides concise, factual briefings intended to support customer Legal, Privacy, Information Security, and Risk stakeholders when evaluating specific aspects of the Preservica platform. Each document in the series focuses on a defined topic and is designed to complement, not replace, formal contractual documentation and due diligence materials.

These briefings are provided for informational and assurance purposes only. They describe Preservica's approach, controls, and practices as of the date of publication and do not constitute legal advice, contractual commitments, or warranties. Customers should rely on their applicable agreements with Preservica, including any data processing agreements and service terms, for binding obligations.

This document forms part of the Preservica CISO Assurance Series and may be read independently or alongside other series publications, depending on customer needs.

Preservica

Preservica

1. About this document

Preservica provides AI-enabled features within its digital preservation platform to help customers discover, describe, and manage large volumes of digital content more efficiently.

This document is designed as a **one-stop assurance pack** that customers can share internally with Legal, Privacy, Information Security, and Compliance stakeholders when evaluating or approving the use of Preservica AI features.

It explains:
- How Preservica uses AI in practice (without marketing hype)
- The security, privacy, and trust controls built into Preservica AI features
- How responsibilities are shared between Preservica, its technology partners, and customers
- Why Preservica AI can be adopted with confidence and minimal internal friction

This document is **informational in nature**. It supports customer understanding and internal assurance discussions and should be read alongside applicable contractual terms, data processing agreements, and regulatory obligations.

2. Preservica AI at a glance

Preservica's AI features are designed to support specific, clearly defined use cases commonly required in archival and digital preservation workflows, such as:

- Optical character recognition (OCR)
- Identification of potential personally identifiable information (PII)
- Image categorization and description
- Audio and video transcription
- Metadata quality checks and standardization

Across all AI features, Preservica applies a consistent set of design principles:

- **Customer choice and control** – AI features are opt-in and must be explicitly enabled. No AI processing occurs by default.
- **Human-in-the-loop** – AI outputs are presented for user review. Preservica AI does not take autonomous action on customer data.
- **Purpose-limited use** – AI is used only to perform the task requested by the user.
- **No silent learning** – Customer data is not used to train, fine-tune, or improve AI models.

3. Responsible AI and regulatory alignment

3.1 Preservica's role under the EU AI Act

Under Regulation (EU) 2024/1689 (the EU AI Act), Preservica acts as a Deployer of AI systems when providing AI features as part of its SaaS platform.

In practice, this means Preservica:

- Selects and integrates AI models for defined, low-risk use cases
- Applies appropriate technical and organizational safeguards
- Ensures relevant staff have appropriate AI literacy and training
- Implements human oversight and usage controls

Customers retain control over:

- Whether AI features are enabled
- Which content is submitted for AI processing
- How AI outputs are used within their organization

3.2 A globally consistent approach

While AI-specific regulation continues to evolve globally, Preservica applies the same core responsible-AI principles across all regions, including:

- Risk-based design
- Transparency of processing
- Data minimisation
- Human oversight

This approach supports alignment with emerging and existing expectations in the UK, United States, Australia, and other jurisdictions.

4. How Preservica deploys AI

4.1 AI operated entirely within Preservica

AI models and processing components are deployed and operated fully within the Preservica SaaS service boundary.

- Customer data remains within Preservica-controlled infrastructure
- No prompts, inputs, or outputs are shared with third-party AI services

Examples include OCR, PII identification, and audio/video transcription using open-source technologies.

4.2 AI services managed by Preservica within hyperscaler environments

Preservica may use AI services provided by the same hyperscale cloud provider that already hosts a customer's Preservica environment (for example, Microsoft Azure or Amazon Web Services).

In this model:

- AI services operate outside the core Preservica SaaS boundary
- Processing remains within the customer-approved cloud environment
- Preservica manages configuration, data flows, and contractual arrangements. Key safeguards include inference-only usage, no model training on customer data, stateless processing, and regional controls.

4.3 Customer-managed AI integrations

Customers may choose to connect Preservica to AI services that they select and manage independently.

- Customers control service selection, configuration, and credentials
- Responsibility for the external AI service rests with the customer

Preservica can provide configuration support through Professional Services where required.

5. AI features and technologies used today

The table below provides a high-level view of selected Preservica AI features and the primary technologies used to deliver them.

AI Feature	Purpose	Primary Technology / Model	Deployment Model
Optical Character Recognition (OCR)	Text extraction from images and scanned documents	Tesseract[1]	Preservica-operated
PII Identification	Detection of potential personal data	Microsoft Presidio[2]	Preservica-operated
Audio Transcription	Speech-to-text conversion	Whisper[3]	Preservica-operated
Video Transcription	Speech-to-text conversion	Whisper[3]	Preservica-operated
Image Description	Automated image categorization and description	Hyperscaler-provided vision models[4]	Preservica-managed

Preservica continuously evaluates AI technologies to ensure they remain appropriate, secure, and effective for customer needs. **Preservica may change or replace underlying AI models or technologies over time**, provided that the security, privacy, and trust controls described in this document are maintained.

6. Protecting customer data

6.1 What data is processed

Depending on the AI feature enabled, processing may involve:

- Customer-selected content
- Preservica-defined prompts required to perform the requested task
- Generated outputs (such as transcriptions or classifications)

6.2 How data is protected

- AI models do not retain customer prompts or outputs
- Customer data is not reused across customers
- Customer data is not used to train or fine-tune AI models

All data is encrypted in transit using TLS 1.2 or higher and remains encrypted at rest in line with Preservica's standard platform controls.

7. AI model training and provenance

Preservica AI features rely on pre-trained models supplied by established technology providers or open-source projects.

- Model training is performed by the original providers prior to Preservica integration
- Preservica does not train or fine-tune models using customer data
- Customer data is never incorporated into training datasets

Public documentation from model providers is available for customers who wish to review training approaches and responsible-AI practices.

8. Security, assurance, and oversight

- AI processing is logically isolated per customer environment
- No cross-tenant data access occurs
- AI services are monitored as part of Preservica's standard security operations
- Human review is required before AI outputs are relied upon operationally

Where applicable, AI services provided by hyperscale partners fall within the scope of their independently assessed certifications (such as ISO 27001 and SOC reports).

9. Evolving with confidence

This document reflects Preservica's current approach to delivering AI features in a secure, responsible, and transparent way.

As technology, regulation, and customer expectations evolve, Preservica may update its AI features and underlying technologies while maintaining appropriate levels of security, privacy, and human oversight.

10. Next steps

Customers with additional legal, privacy, or security questions are encouraged to raise them through their usual Preservica contact or formal due-diligence process.

Preservica is committed to helping customers adopt AI features with confidence in support of long-term digital preservation.

11. Frequently Asked Questions (FAQ)

The following FAQs are based on common questions we receive from customer Legal, Privacy, Information Security, and Compliance teams (including customer questionnaires and due-diligence follow-ups).

11.1 What AI features does Preservica provide today?

Preservica provides AI-enabled capabilities designed for practical archival and digital preservation workflows (for example OCR, PII identification, transcription, image description, and metadata quality support). The specific feature set available to a customer depends on the Preservica product/edition and configuration.

11.2 Are Preservica AI features enabled by default?

No. Preservica AI features are **opt-in** and must be explicitly enabled before any AI processing occurs. AI outputs are presented to end users for review and acceptance.

11.3 Is there “human-in-the-loop” oversight?

Yes. Preservica is designed so that AI outputs are provided as suggested results for user review rather than autonomous actions taken on customer data.

11.4 Do you use customer data to train or fine-tune AI models?

No. Preservica does not use customer data to train or fine-tune AI models. Where Preservica uses third-party models or services, they are used for inference only.

11.5 Where does the model training data come from (if you don’t train on customer data)?

The AI models used by Preservica features are **pre-trained** by their respective providers or open-source projects prior to Preservica integration. Preservica does not control or supplement those training datasets. Customers who require deeper detail should refer to the providers’ published documentation referenced in this document.

11.6 Do prompts, inputs, or outputs get retained by the AI model or service?

No. Prompts, inputs, and outputs generated when using Preservica AI features are not retained by the AI models for training or learning purposes.

11.7 Are you using any third-party AI providers?

Depending on the feature and deployment model, Preservica may use open-source components and/or hyperscaler AI services for inference. Preservica also supports customer-managed integrations where customers select and manage the external AI provider themselves.

11.8 How is customer data protected when AI features are used?

Preservica applies its standard platform security controls (including encryption in transit and at rest) and designs AI features to be purpose-limited, opt-in, and subject to human review. For externally deployed hyperscaler AI services, the relevant provider security and data-protection commitments also apply.

11.9 How do you handle prompt injection and similar AI-specific risks?

Preservica reduces exposure to AI-specific risks by limiting AI features to clearly defined use cases, maintaining human oversight, and applying appropriate security and configuration controls when integrating AI capabilities.

11.10 Do you log AI usage for audit and operational assurance?

Preservica maintains standard operational logging and monitoring consistent with its platform security practices. Where customers require specific audit evidence relating to AI usage, this can be addressed through the normal due-diligence and assurance process.

11.11 Can we restrict which users can access AI features?

Yes. AI features are enabled through administrative control and standard role-based access mechanisms, allowing customers to restrict access to authorized users and defined use cases.

11.12 How will we be informed when AI technologies or models change?

Preservica communicates material product and capability changes through its normal customer communication channels (for example release notes), and AI features remain subject to customer opt-in control.

11.13 Does Preservica train its staff on AI-related topics?

Yes. Preservica ensures that relevant staff receive appropriate training aligned with their roles, including AI awareness and literacy, responsible use, information security, and data protection/privacy obligations. This training forms part of Preservica's broader security and compliance programme.

11.14 How does Preservica provide assurance around AI security, reliability, and compliance?

Assurance for Preservica AI features is provided through Preservica's existing information security and compliance programme. This includes independently assessed certifications, audits, and customer-available assurance reports (for example ISO 27001 and SOC reports), which cover the underlying platform controls on which AI features operate.

Appendix A - Reference materials

This appendix provides links to publicly available documentation published by technology providers referenced in this document. These materials are provided for customer convenience and independent review.

A.1 Open-source technologies

[1] Tesseract OCR – Project documentation, architecture, and training information:
https://tesseract-ocr.github.io/tessdoc/

[3] OpenAI Whisper – Open-source speech-to-text model documentation and model cards:
https://github.com/openai/whisper

A.2 PII detection and privacy tooling

[2] Microsoft Presidio – Data protection, privacy design, and governance documentation
https://microsoft.github.io/presidio/

A.3 Hyperscaler AI governance and data privacy

[4] Microsoft Azure Responsible AI – Responsible AI principles, governance, and data privacy practices:
https://azure.microsoft.com/en-us/solutions/ai/responsible-ai-with-azure/

[4] AWS Responsible AI and data governance – Security, privacy, and governance considerations for AI services:
https://docs.aws.amazon.com/prescriptive-guidance/latest/strategy-data-considerations-gen-ai/security.html

Preservica may reference alternative or additional provider documentation where appropriate as underlying technologies evolve. Inclusion of reference materials in this appendix does not imply endorsement of any specific vendor beyond their use within the Preservica platform.

About Preservica

Preservica is transforming the way organizations around the world protect and future-proof critical long-term digital information. Available in the cloud (SaaS) or on-premise, our award-winning Active Digital Preservation™ archiving software has been designed from the ground up to tackle the unique challenges of ensuring digital information remains accessible and trustworthy over decades.

It's a proven solution that's trusted by thousands of businesses, archives, libraries, museums and government organizations around the world, including the UK National Archives, Texas State Library and Archives, MoMA, Yale and HSBC.

preservica.com/about