



AI Use Policy

Pocket Museum · Cloud and Culture Ltd · Effective 26 August 2026

Last updated: 26 August 2026. Owner: the Director of Cloud and Culture Ltd. Contact: support@cloudandculture.co.

1. Why we are publishing this

Pocket Museum hosts story-led cultural experiences, meaning virtual museums, walking tours and curated exhibitions, created by institutions, communities and individual curators. Trust in what we publish is the product. Visitors need to know where a story came from, whether an image depicts something real, and who stands behind the claims on the page.

We use artificial intelligence in parts of how we build and operate the platform. This policy states plainly where we use it, where we do not, and what holds regardless.

This policy applies to Cloud and Culture staff, contractors, and anyone publishing through the platform. It sits alongside our Terms and Conditions, our Privacy Policy and our Community Standards. Where those documents and this one differ, the Terms govern.

2. Our commitments

These hold across every use of AI described below.

- **A named person is accountable for everything we publish.** No AI-generated or AI-assisted museum, exhibition, story or label reaches a visitor without review and sign-off by a named person.
- **Real over synthetic.** For exhibit imagery we prioritise authentic photographs, scans and documentation of real objects, places and people over AI-generated visuals.

- **We disclose.** Where AI-generated media appears in an exhibit, it is labelled where the visitor meets it, not only on a policy page. Synthesised narration is disclosed in our Terms and Conditions.
- **Automated decisions about visitor comments are limited, disclosed and reversible by a person.** An automated tool does not hide a comment that it judges to break our Community Standards before any person has seen it. It never restricts or suspends an account, it never decides an appeal, and you can always see your own comment together with the decision we reached. A member of staff reviews what the automated tools have decided at least once a month.
- **Cultural material gets extra care.** AI assistance is not a substitute for consulting source communities, rights holders and originating institutions.

3. Where we use AI

3.1 Narrative development and storytelling assistance

We use generative AI to help draft, structure and refine exhibit narratives, tour scripts, object labels and supporting copy, for example proposing outlines, suggesting alternative framings for a story, tightening prose, adapting reading level, or drafting variations for different audiences.

Constraints:

- AI output is a draft, never a final text.
- Curators and editors verify every factual claim against primary sources before publication.
- Attributed quotes, oral histories and community testimony are never AI-generated or AI-paraphrased. They are reproduced from the source record.
- Where a partner institution supplies its own scholarship, AI is not used to rewrite that scholarship without the partner's agreement.

3.2 Fact-checking and research support

We use AI tools to cross-check dates, names, provenance details and factual assertions in draft content, and to surface possible errors, gaps or unsupported claims for human follow-up.

Constraints:

- AI flags are a prompt to check a source, not a source themselves.

- Verification is completed against primary or authoritative secondary sources by a person.
- A tool's failure to flag something is not treated as confirmation that it is correct.

3.3 Checking drafts against our own publication standards

Before a museum is published we run its draft text through an automated review against our published content criteria, covering editorial standards, accessibility requirements, rights and licensing rules and sensitivity guidance. It flags things like a missing image credit, absent alternative text, unclear provenance, or language that may need consultation with a source community. Each finding is recorded against the criterion it relates to, with a severity, for an editor to act on.

Constraints:

- These checks assist our editorial team. They do not replace editorial judgement.
- A clean automated check is never sufficient grounds to publish on its own.
- The review reads the draft museum only. It does not read anything about the visitors who will see it.

3.4 First-line moderation of visitor comments

Every comment and reply posted to Pocket Museum is screened automatically against our Community Standards before it is shown to other readers. Messages sent to us through the contact form are screened the same way, to separate genuine enquiries from spam and abuse.

How it works, stated as it actually runs:

- **An automated tool checks every comment first.** Where it judges the comment to clearly break our Community Standards, the comment stops being shown to other readers straight away, before a person has looked at it.
- **Where the tool is unsure, a person decides instead.** Ambiguous context, a claim that might be defamatory, disguised swearing and borderline vulgarity are all held for human review rather than refused.
- **We aim to complete that human review within two working days.**
- **We never restrict or suspend an account by automated means.** Only a person does that, and only after warning, except for serious breaches.

- **You can always see your own comment**, whatever we have decided about it, marked with the decision we reached and why. Where automated tools were used to reach a decision about your content, we tell you so.
- **A member of our staff reviews what the automated tools have decided at least once a month**, and we correct anything we find to be wrong.
- **Appeals are decided by a person, never by a tool.** You can challenge any decision free of charge for six months from the day we tell you about it, by emailing support@cloudandculture.co with Appeal in the subject line. If an appeal shows our decision was wrong we reverse it without undue delay and tell you.

What is sent to the model is the text of the comment or message. Your name, email address, account identifier and browsing history are not sent with it, and the model is not given any record of you.

We monitor moderation outcomes for error and bias, including differential effects on non-English submissions, minority-language content and culturally specific material, and we adjust our thresholds and guidance accordingly. That monitoring happens as part of the same monthly review.

3.5 Synthesised narration

Spoken-word narration in a museum is often synthesised rather than recorded, using a text-to-speech model. Only the museum's own written text is sent for that, never anything belonging to a visitor. The audio is generated once by the contributor and stored with the museum.

Constraints:

- Narration text forms part of the museum and is reviewed by us before publication like any other content. Review does not make generated speech accurate, and our Terms and Conditions say so.
- We do not synthesise the voice of an identifiable real person.
- Where a recording of a real voice exists and may be used, we prefer the recording.

4. Images and visual media

Pocket Museum is in the business of showing people real things. Our default is documentary imagery: photographs, scans, 3D captures and archival reproductions of actual objects, sites and events, properly credited. No AI-generated imagery is published in Pocket Museum today, and the platform has no image generation feature. The rules below govern any use we introduce.

We use AI-generated or AI-enhanced imagery only where all of the following hold:

- No authentic visual record exists or is available under an appropriate licence, and
- the image serves a clear interpretive purpose, such as a reconstruction, a diagram, or an atmospheric background clearly distinct from the historical record, and
- it is visibly labelled as AI-generated where it is displayed, and
- it has been approved by the responsible curator, and by the partner institution where the exhibit is theirs.

We do not:

- Present AI-generated images as authentic historical photographs, documents or artefacts.
- Generate synthetic depictions of identifiable real people, living or deceased, as if documentary.
- Use AI to fabricate or complete missing parts of an artefact, artwork or archival image in a way that could be mistaken for the original. Reconstructions are shown alongside the original and labelled.
- Use AI upscaling or restoration on archival material without noting that the reproduction has been digitally processed.

We retain provenance metadata for the images we publish, including whether a given asset is documentary or synthetic.

5. Where we do not use AI

- To decide an appeal, or to restrict, suspend or close an account.
- To make the final decision on whether a museum, exhibition or story is published.
- To generate content presented as the words of a named person, community or institution.
- To generate scholarship, provenance claims or attributions that are then published without expert verification.
- To replace consultation with source communities and rights holders on sensitive cultural material, including material subject to Indigenous data governance principles, restricted knowledge or repatriation claims.

- To profile visitors, to target advertising, or to personalise what a child sees. We show no advertising at all.
- To make decisions about hiring, pricing, or whether to work with a partner institution.

6. Cultural sensitivity and community control

Cultural heritage material carries obligations that general-purpose AI tools do not understand. Some knowledge is restricted, some images should not be reproduced, and some stories belong to communities who decide how they are told.

- Where a source community or originating institution has placed conditions on how material may be used, those conditions govern, including whether AI tools may be applied to it at all.
- Communities contributing to Pocket Museum may ask that their material be excluded from AI-assisted processing. Write to support@cloudandculture.co.
- Where our automated checks flag potentially sensitive material, the matter goes to a person with relevant expertise, and to the partner where appropriate.

7. Data handling

- **Anthropic** provides the model behind comment moderation, contact-form screening and the pre-publication content review, through the Claude API. Anthropic acts as our service provider under its commercial terms and data processing addendum, which includes the UK Addendum.
- **Google** provides the text-to-speech model used for synthesised narration. Only museum text is sent to it.
- **Content submitted to Pocket Museum by visitors and partners is not used to train third-party AI models.** Our agreements with these providers do not permit our inputs to be used to train their models.
- Personal data is processed in accordance with our Privacy Policy, which describes what reaches an AI model, on what legal basis, and for how long the result is kept. Our data is stored in the United Kingdom. Processing by Anthropic takes place outside it, under the transfer safeguards set out in that policy.
- Partner institutions can ask for our specific commitments on collection data handling at support@cloudandculture.co.

8. Accountability

- **Responsible owner.** The Director of Cloud and Culture Ltd holds accountability for AI use across the platform.
- **Editorial sign-off.** Every published museum records the person who approved it.
- **Review cadence.** This policy and our AI tooling are reviewed at least every twelve months, and whenever we adopt a materially new use.
- **Incidents.** If AI-assisted content results in a published error, we correct it, note the correction on the page, and record the cause in our internal log.
- **Questions, concerns, or reports of AI-related errors.** Write to support@cloudandculture.co.

9. Changes to this policy

We will update this page when our practices change, and note the date of the change above. Material changes will be communicated to partner institutions directly.