

CULTURE, RELIGION, VALUES AND HISTORY —
THOUGHTFULLY EXPLORED

TheCulturePost

A workspace for well-sourced cultural
storytelling, citable heritage archives,
and dialogue across traditions that
stays generous.

CITATION-FIRST EDITORIAL

PROVENANCE-AWARE ARCHIVE

ORAL HISTORIES

MODERATED DIALOGUE

CITATION-GROUNDED AI



01 The short version

TheCulturePost is one place to read about culture, faith and history where every claim shows you the source behind it — and to preserve the fragile knowledge that would otherwise quietly disappear.

THE PROMISE

If a claim cannot be traced to a source, it does not get published. Everything else in the product exists to make that practical rather than aspirational.



Every topic, story, citation, collection, archive item, oral history and discussion belongs to a workspace — which is what lets the same product serve a curious reader and a national archive with fifty staff.

Broad content with no sources, or academic material behind a paywall

Almost everything written about culture and religion online sits at one of two extremes. The casual reader gets noise. The serious reader gets shut out. The communities whose traditions are being described rarely get a say in how they are represented, and the most fragile knowledge of all — the part that lives only in an elder's memory — keeps slipping away while everyone argues about the rest.

TheCulturePost sits in the middle and does three things at once: a citation-first editor where a footnote stays attached to the claim it supports; a heritage archive built for provenance, where a scan, photograph or recorded testimony carries its origin, rights and access level; and a place to talk about all of it, with visible norms, pinned context and moderation on every post.

01 What holds it together

Every claim

Traceable to a source

A story carries its citations, each with source type, author, publisher, excerpt and a chain-of-custody note.

4 tiers

Rights on every item

Public, community, restricted and embargoed — set per archive item and per oral history, not per site.

6 roles

Seeded, not configured

Editor-in-chief through community member arrive in every space at creation. Nobody builds a permission matrix by hand.

Pre-launch

Waitlist open for 2026

No users, no revenue and no track record are claimed. This document describes what is built and what is not.

02 Who it is for — readers and makers

Curious readers

They want to understand a festival they were invited to, a sacred text, or a turning point in history — in a few quiet minutes, with context and sources rather than a search-results argument.

Educators and students

Teaching and learning that rests on primary sources anyone can check. Classroom and curriculum tooling is on the roadmap; curated collections and topic hubs are available now.

Contributors and writers


Essayists, journalists, scholars and community storytellers who want to publish responsibly and are tired of tools that treat sourcing as an optional plugin.

Diaspora and community organisers

People preserving the traditions they carry and convening exchange across faiths, who need a space designed for generosity rather than one moderated after the damage.

Editors and platform stewards

The people who run fact and sensitivity review, approve contributors, and are accountable when a piece about a living tradition gets it wrong.




Who it is for — custodians and institutions


Archivists and heritage custodians

Communities and individuals capturing oral testimony and describing primary sources, who need provenance, rights and access levels as fields rather than as a spreadsheet convention.

Cultural and heritage institutions

Museums, libraries, universities, schools and faith organisations that want a branded, governed home for their own collections rather than a consumer blog.





03 What breaks today

None of these problems are new, and none of them are solved by writing better content. They are structural — they come from tools that make the careless thing easy and the careful thing expensive.

03 Sourcing, memory and hostile threads

The confident, unsourced claim

Most cultural and religious content online flattens a tradition or repeats a misconception because nobody checked, and there is no way for a reader to tell the difference.

- Sourcing is structural rather than editorial etiquette. A story holds its citations as first-class records, and the review gate stands between a draft and publication.

Fragile knowledge is disappearing faster than it is captured

An elder's testimony, a single manuscript, a recording on a failing drive — the most precious material is the least durable, and it is nobody's job to save it.

- Oral-history capture and a provenance-aware archive turn perishable memory into a described, citable, findable record with attribution and access rules attached.

Conversation across traditions turns hostile by default

The threads that could build understanding become the ones nobody wants to host, so institutions disable comments and the exchange never happens.

- Discussions are a modelled primitive: anchored to a story or a question, framed with pinned context, markable as an interfaith space, with a moderation state on every reply.

03 Sacred material, governance and evidence

Sensitive and sacred material has nowhere safe to live

Ordinary publishing tools offer public or private and nothing in between, so custodians either over-expose material or keep it offline entirely.

- Four access levels — public, community, restricted and embargoed — sit on every archive item and oral history, and restricted material is reached through a separate elevated-access path.

Institutions get a blog when they need a governed workspace

A museum or faith organisation is handed a consumer publishing tool with no roles, no audit trail, no single sign-on and no way to separate the archivist from the writer.

- Branded spaces come with a seeded role catalogue, single sign-on, an audit and activity timeline, quotas and billing — the same governance a procurement team expects.

Readers cannot tell verified from contested

Two claims that look identical on the page can rest on a peer-reviewed source and on a travel blog, and the reader is left to guess.

- A typed source taxonomy across eight kinds and an editorial verification status — unverified, corroborated, disputed, retracted — make the strength of the evidence visible on the claim itself.

03 AI and fragmentation made it worse

AI has made the sourcing problem considerably worse

A model will happily explain a festival's origin in fluent prose, and there is no way to tell whether that came from a scholar, a listicle or the model's own pattern-matching.

- Generated answers are built only from the sources attached to the record you asked about, return the exact list they were given, and are refused outright when the record carries no citations.

Cultural knowledge is fragmented across languages and walled gardens

The same tradition is documented in four languages on five platforms, none of which can cite the others, and none of which a researcher can query.

- Topic areas, timelines and collections gather material into one discoverable spine, and public holdings are exposed through a read API and a standard image endpoint that heritage viewers already speak.

04 What you get

Nine building blocks make up the product. The first six are the domain itself; the last three are what turns a publishing tool into something an institution can sign off on.

DISCOVERY

A library you can actually narrow

Filter essays, explainers, photo-stories, timelines and oral histories by tradition, region, era, theme and topic area. Reading lists carry across devices.

EDITORIAL

A citation-first editor and a real lifecycle

Inline citations, source attachments and footnotes tied to the claims they support. Every story moves through idea, draft, in-review, published and archived.

EVIDENCE

Citations as records, not footnotes

Each citation carries a source type from eight kinds, author, publisher, link, excerpt, publication date and a chain-of-custody note.

04 Verification, archive, testimony

VERIFICATION

Editorial confidence on the claim

During review an editor marks a cited claim unverified, corroborated, disputed or retracted. The status travels with the claim, so a reader sees the confidence behind an assertion rather than only the assertion.

ARCHIVE

Heritage material with provenance and rights

Seven item kinds — manuscript, document, photograph, audio, video, object and map — each with provenance, origin, descriptive metadata and one of four access levels, grouped into collections curated by tradition.

TESTIMONY

Oral histories captured before they fade

A named narrator, a transcript, language, duration and an access level, stored durably with the transcript encrypted at rest. Media is attached by upload today; recording inside the app is planned.

04 Community, intelligence, governance

COMMUNITY

Dialogue designed to stay generous

Threaded discussions anchored to a story or a question, framed with a pinned prompt, markable as an interfaith space, with perspective tags on replies, four moderation states and lock and reopen controls.

INTELLIGENCE

AI that answers from your sources or not at all

Four grounded tasks run against your own citations, holdings and testimony, each returning the exact source list it was handed, declaring its own uncertainty, and landing as a pending proposal for an editor.

GOVERNANCE

The part procurement asks about

Six seeded workspace roles across eight resource types, permission checks at the API edge, workspace isolation on every write, an audit and activity timeline, quotas enforced before the action, and billing.

05 Put to work — sixteen scenarios

TheCulturePost has not launched, so none of these describes a real customer — each describes what the platform is built to do, with the parts that are not built yet named as such.

CURIOUS READER

Understand a tradition before you walk into it

CURIOUS READER

Walk a tradition across centuries instead of reading dates

EDUCATOR

Teach from primary sources rather than summaries of summaries

CONTRIBUTOR

Write a piece where the sourcing survives the edit

MANAGING EDITOR

Run review as a gate rather than a courtesy

REVIEWING EDITOR

Mark how strong the evidence actually is

HERITAGE ARCHIVIST

Capture an elder's testimony before there is nothing to capture

HERITAGE ARCHIVIST

Turn a shelf of primary sources into something citable

COMMUNITY MEMBER

Discuss a difficult subject without the thread collapsing

COMMUNITY ORGANISER

Bring traditions into conversation on equal footing

DIASPORA COMMUNITY

Find the people who carry the tradition you are reading about

EDITORIAL STEWARD

Keep a multi-tradition publication healthy as it grows

INSTITUTION

Run a branded, governed space for a collection you steward

MANAGING EDITOR

Summarise the evidence behind a story, from that evidence only

CURATOR

Find the threads running through your own collection

PLATFORM STEWARD


Run the editorial desk from the AI client you already work in


Capture an elder's testimony before there is nothing to capture

The bind. The knowledge most worth preserving lives in a person's memory, and there is no second chance to record it.

You record the testimony, transcribe it, and describe it: the narrator, the language, the duration, the tradition and the place. You set an access level appropriate to what was shared and what was consented to. The transcript is held encrypted at rest, and the record joins a collection where contributors can cite it and readers can find it.

- A single voice becomes part of a durable record. Recording inside the app is planned; today the media is attached by upload.





MANAGING EDITOR

Summarise the evidence behind a story, from that evidence only

The bind. A model that has read the whole internet will explain a festival's origin beautifully, and you cannot tell whether that came from a scholar or from nowhere.

You ask what the sources behind a story actually say. The service gathers the citations attached to that story — up to twenty-five, rendered as a numbered block with title, author, publisher, source type and excerpt — and instructs the model to answer from those alone, cite them inline, and say the sources are insufficient rather than fill the gap. A story with no citations is refused outright.


-
- Reachable through the API and connected agents today; there is no button for it on a story yet.

06 Built AI-native

AI that answers only from your citations — and an editor still signs it off

In a setting that covers faith, heritage and contested history, the failure mode of AI is not an ungrammatical sentence; it is a fluent, confident sentence with nothing behind it. So the product does not try to make a general model behave. It removes open-ended recall from the set of available options.

Everything starts from grounding: the sources are gathered from your own space, that list is fixed at that moment and stored with the run. A story with no citations produces a refusal, not a helpful guess.




06 Four rules the AI layer runs on




Your sources go in first

Answers are assembled from the citations, archive items and testimony attached to the record you asked about. An empty source set is a refusal rather than an improvisation.



It shows exactly what it read

Each answer comes back with the precise list it was handed — captured before the answer exists, so grounding is something you verify rather than take on trust.



Nothing counts until an editor says so

Answers arrive pending. Approving or rejecting is one recorded decision, taken once, and nothing is written back into your story on your behalf.



Sealed material is never sent to a model

Only public and community-tier holdings are ever used for grounding, even if the person asking has every right to open the rest themselves.

06 What the AI can actually do — the four grounded tasks

Evidence summary from a story's own citations

Up to twenty-five citations are rendered as a numbered block with title, author, publisher, source type and excerpt.

Answer only from these, cite them inline, and say they are insufficient rather than fill the gap.

Related traditions across a collection

Ask what connects the holdings — shared motifs, migration routes, an observance under three local names.

It reasons across metadata from that collection alone, so every suggested connection walks back to specific items.

Respectful framing draft

Draft a framing note for a piece about a living tradition, built from the material attached to that record.

Framing is where cultural writing most often goes wrong, and a grounded draft gives an editor something to argue with.

Oral-history transcript cleanup

Tidy the paragraphing and punctuation of a transcript a person already typed. This is explicitly not speech recognition, and no audio is ever sent to a provider.

The captured transcript stays the record of what was said; the tidied version has to be approved.

06 How every answer is constrained

Refusal on an empty source set

A story with no citations cannot be summarised, and a collection with no eligible holdings cannot be analysed. The request fails with a message telling you what to add.

Self-declared uncertainty

Each answer opens with a low, medium or high rating, parsed out and stored alongside it. A missing rating is filed as high rather than silently as low.

A source list you can check

Every answer returns the exact citations, archive items and oral histories it was given, each one openable, captured before the answer exists.

Pending-by-default human review

Nothing auto-approves. There is no configuration in which a generated response publishes itself, and a failed run can never acquire an editor's signature.

Access-tier floor on grounding

Grounding reads only public and community-tier material. The limit lives inside the AI layer, so relabelling access levels later cannot quietly widen what leaves the building.

One decision per run

A second attempt on an already-reviewed run is rejected. Two editors acting together produce one winner and a clear error, not a silently replaced judgement.

06 Ledger, permissions and reach

Append-only run ledger

One entry per request, success or failure: who asked, what it was grounded on, which model answered, the answer, its uncertainty, timing and the review outcome.

Separate AI permissions

Viewing run history, generating an answer and reviewing one are three distinct grants. Three of the six role templates carry any AI rights at all.

Context-aware assistant on every screen

A floating assistant resolves the story, topic, discussion, collection or archive item you have open from the route you are on. The payload that leaves the browser is a fixed list of nine values.

Thirty-one tools for external agents

The whole editorial, archive and community domain over the Model Context Protocol, with twenty-one shared workspace modules alongside it, acting under your own permissions.

Five model choices with a per-request override

Gemini 2.5 Flash as the default, plus Gemini 2.5 Pro, GPT-4o mini, Claude 3.5 Haiku and auto-routing through OpenRouter.

Reader explanations and translation

Planned: tap a term for a plain explanation tied to a citation, and read a piece in your own language with the original one tap away. Designed and sequenced, not built.

06 And what keeps it in bounds

Grounding you can audit

The deterministic list supplied to the model is stored with the run and openable item by item. It is never reconstructed from what the model afterwards claims it used.

A floor, not a permission check

A curator who can personally open a sealed testimony still cannot cause it to be forwarded to a provider.

Human decision on the record

Approval or rejection records the actor, the timestamp and optional notes, can be taken once, and is impossible on a run that produced nothing.

Non-identifying request summaries

Raw prompts, reader-supplied text, story bodies, replies and citation excerpts are not stored in the log.

Testimony decrypted for one call only

Transcripts are unlocked for the duration of a single request, and only the title and transcript enter the prompt — never the narrator record.

AI can be switched off independently

A separate feature switch lets an operator disable AI generation without disabling the rest of the product.

06 Where the AI is reachable, and in what numbers

REACHABLE FROM

- The floating assistant, on Home, the editorial overview, Stories, Topics, Timeline, Discussions, Archive and Saved
- The governed API, where the four citation-grounded tasks run today
- A Model Context Protocol server with thirty-one domain tools, twenty-one shared modules, four read-only resources and two guided recipes
- The story screen panel and the editorial review queue — designed and ticketed, not built
- Reader-facing explanation, translation and recommendation — designed, not built

IN NUMBERS

- Four grounded tasks run today: evidence summary, framing draft, related traditions across a collection, and transcript cleanup
- Up to twenty-five citations are rendered into the numbered source block for an evidence summary
- Three separate AI permissions — view history, generate, review — and three of six role templates carry any of them
- The assistant sends a fixed list of nine context values and no story text, replies, transcripts or citation excerpts
- No accuracy or quality figure is claimed anywhere, because no evaluation set for the AI path exists yet

06 What is not there, plainly

The citation-grounded generator, the review gate and the run ledger all work, but no screen in the app calls them — there is no button on a story and no queue on the editorial overview, so today they are reached through the API and connected agents.

The assistant is a general platform agent rather than a TheCulturePost specialist, and it is not wired to the grounded path. Agent actions that change state go through on first ask; the confirmation gate exists in the shared agent core and is not switched on here.

Cost per workspace is not calculated and there is no evaluation set, so no spend, accuracy or quality number is quoted. Bring-your-own-key is not supported. Audio speech-to-text does not exist. The entire reader-facing layer is designed and sequenced, not built.

07 Who does what — editorial roles

Six roles are seeded into every space at creation, plus organisation and billing presets. In several of these roles what is withheld is the point.

Editor-in-Chief

The publication steward — curates topic areas, oversees content across traditions, and holds the only role with full authority over all eight resource types.

Cares about editorial integrity across the whole publication, and the ability to withdraw anything.

Managing Editor

Owns the gate between draft and published: fact and sensitivity review, verification of cited claims, revision requests, and topic curation.

Cares whether a piece holds up before it goes live, and whether its evidence is honestly labelled.

Contributor

The writer — drafts essays, explainers, photo-stories and timelines in the citation-first editor and attaches the sources that back every claim.

No authority over the archive. Reads archive items and oral histories only to cite them.

07 Custodians, moderators, members and owners

Heritage Archivist

The museum, library or community custodian who captures testimony, describes primary sources with provenance and rights, and owns the archive outright.

Sets the access level on an item, curates collections by tradition, withdraws material on request. No role in moderating dialogue.

Community Member

The core audience seat — reads published stories, follows a claim back to the source behind it, and browses the open heritage archive.

Being able to check anything, without needing to be trusted with anything. No access to oral histories.

Dialogue Moderator

The interfaith convener and space-keeper who opens and frames cross-tradition threads and owns the moderation surface.

Pins context, approves, flags or removes replies, locks and reopens threads. No authority over the archive.

Publication owner and billing manager

Organisation and billing presets generated in the product's own vocabulary — publication owner, editorial admin, editorial viewer, billing manager and billing viewer.

Holds the plan and payment method, reviews invoices and usage, capped below role-assignment escalation.

08 Under the hood

API-first

Product state is served over a federated GraphQL API, so every client — the web app, the assistant, an external AI agent, a partner integration — reaches exactly the same governed surface. There is no privileged back channel and no client-only capability.

Authorisation at the edge

Every query and mutation declares the permission it needs, the gateway checks it before the request reaches the product, and a failed check never runs the resolver. The product itself never validates a token.

Structural isolation

Every write is matched on both the record identifier and the workspace, so an identifier guessed or hallucinated from another tenant returns not-found. Oral-history transcripts are encrypted at rest with authenticated encryption.

THE SPLIT THAT MATTERS

The product owns the citable record; shared workspace storage owns the bytes. An archive item is the described, findable, provenance-bearing record — the manuscript scan or recording behind it lives in shared file storage and is reached through signed URLs.

08 Capability by area

Editorial domain

Topics, stories and citations, with a fixed lifecycle — idea, draft, in review, published, archived — where only valid transitions are allowed, across six story formats.

Citation-grounded AI

Four grounded tasks, a deterministic stored source list, an access-tier floor, self-declared uncertainty, a pending-by-default review gate, and an append-only run ledger.

Evidence model

Eight source types, provenance and chain-of-custody notes, excerpts, an optional scan of the primary source, and a four-value verification status assigned during review.

Access control

Eight product resource types with create, read, update and delete semantics; six seeded workspace roles at distinct priorities; separate organisation and billing presets.

Heritage archive

Seven item kinds with provenance, origin and descriptive metadata; four access levels; collections curated by tradition; oral histories with narrator, transcript, language and duration.

Public read surface

An unauthenticated GraphQL read API and a Level-0 IIIF Image API endpoint for public-access archive items and oral histories, with a compliant capability document.

Community layer

Discussions anchored to a story or question, interfaith marking, pinned prompts, threaded replies with perspective tags, four moderation states, and lock and reopen controls.

Quotas and metering

Allowances on stories, topics, seats, collections, archive items, oral histories, citations, discussions, storage and API requests, enforced before the action.

08 Standards it speaks, systems it plugs into

OPEN STANDARDS IT SPEAKS

- GraphQL, with a federated schema and a public unauthenticated read API
- IIIF Image API at Level 0, so heritage viewers can render public material directly
- Model Context Protocol, for tools, resources and prompts consumed by AI clients
- OpenID Connect single sign-on, enforced on the Enterprise tier
- SAML for institutions that require it, alongside OIDC
- SCIM provisioning for automated joiner, mover and leaver flows
- OAuth 2.0 client credentials and device-code flows for automation and terminal clients
- Outbound webhooks for delivery into systems you already run
- Unicode and right-to-left script support with locale-aware date and number formatting
- Accessibility built to keyboard, screen-reader, focus-management and high-contrast expectations

PLUGS INTO

- Shared file storage for manuscripts, photographs, recordings and cover images, with signed-URL access
- Workspace-wide tagging shared across stories, topics and heritage items
- Inline comments and review threads on stories, citations and archive items
- In-app, email, push and webhook notifications with per-user preferences
- Workspace newsletter and announcements for releases, new collections and editorial news
- Subscriptions, regional pricing, usage-based quotas, add-on capacity and invoices
- A searchable, exportable activity timeline recording actor, target and result
- Calendar and scheduling for publication windows, reminders and recurring operations
- Multi-step workflow orchestration for editorial review, sensitivity gates and access requests
- A developer portal issuing API tokens, OAuth apps and developer keys

08 How it deploys

Multi-tenant cloud

The default. Strict per-workspace boundaries, shared platform services, and no infrastructure for you to run.

Dedicated tenancy

Dedicated database and compute, single-region multi-AZ with private networking and a custom domain, or multi-region with active disaster recovery.

Self-hosted

Run inside your own estate, sized by contributor count, for institutions whose data-handling policy does not permit shared tenancy.

Air-gapped

Fully offline deployment with media-based updates, for national archives and bodies operating on disconnected networks.

Data residency

Regional hosting options so material stays in the jurisdiction that governs it, quoted per engagement.



09 Security & trust — access and isolation

The controls that matter for material that is sensitive, sacred, or attached to a named living person — described as they are built, not as they are aspired to.

Authorisation at the API edge

Every query and mutation declares the permission it requires. The gateway checks it before the request reaches the product, and the resolver never runs on a failed check.

A separate tier for restricted heritage material

Restricted and embargoed holdings are reached through dedicated elevated-access resource types. An ordinary read grant does not reach them.

Tenant isolation by construction

Every write is matched on the record identifier and the workspace together, so an identifier guessed or copied from another tenant returns not-found rather than mutating anything.

Encrypted testimony

Oral-history transcripts are encrypted at rest with authenticated encryption and decrypted only for the duration of a single request that legitimately needs them.

Rights-aware access levels

Public, community, restricted and embargoed sit on every archive item and every oral history, so cultural restrictions and embargoes can be honoured rather than worked around.

Sensitive resources withheld at the role layer

The broad-membership community seat is granted no access to oral histories at all. Only vetted contributors read them, and only archivists and editorial leads can write or withdraw one.

09 Security & trust — review, audit and identity

Editorial verification on cited claims

Unverified, corroborated, disputed or retracted, assigned during review, so the strength of the evidence behind a claim is visible rather than inferred.

Mandatory human review of generated output

Generated content is pending by default, leaves that state only through a single recorded decision, and cannot be approved if the run produced nothing. No setting disables the gate.

Moderation state on every post

Discussions and replies carry pending, approved, flagged or removed, and threads can be locked and reopened, so respectful dialogue is enforceable.

Audit and activity timeline

Every action is recorded with actor, target and result, searchable and exportable, with retention scaling by plan from seven days to a full year and beyond.

The AI access-tier floor

Restricted and embargoed material is never sent to a model, whatever rights the person asking holds. The rule lives in the AI layer itself.

Quotas enforced before the action

Allowances warn at eighty per cent, block new usage at one hundred, and never auto-purchase additional capacity, so there is no after-the-fact bill surprise.

Append-only AI run ledger

One entry per request, changeable only in its review outcome and isolated per workspace at the row level. Requests are stored as short non-identifying summaries.

Identity, encryption and disclosure

Single sign-on with MFA, encryption in transit and at rest, enforced SSO with SAML or OIDC and SCIM on Enterprise, and a published responsible-disclosure route.

STATED PLAINLY

TheCulturePost holds no security certifications today

It is pre-launch, and the security practices described are modelled on the ISO 27001 framework and designed against SOC 2 criteria — SOC 2, ISO 27001 and ISO 27701 sit on a certification roadmap, not in a filing cabinet.

Data-protection alignment with GDPR and India's DPDP regime is by design, and a data processing agreement is available, but nothing here should be read as an audited attestation. Compliance evidence exports and SIEM streaming are on the roadmap.

09 Three known gaps in the permission model

Worth stating before you buy rather than after. Separate publish, verify and moderate actions are the fix, and they are not shipped.

PUBLISH

Advancing a story through its lifecycle is enforced as the same action as editing its body, so any role that can draft can also publish — which weakens the review gate that is the product's central promise.

VERIFY

Setting a citation's verification status is enforced as the same action as correcting its title, so a contributor can in principle mark their own source corroborated.

MODERATE

Moderating a reply is the only operation behind the reply update action, which is why the broad community seat is held read-only rather than being given posting rights that would also carry removal power over everyone else.

Community consent, takedown and access-request workflows are specified as first-class auditable flows and are not yet built.

10 One product, many front doors

Some of these front doors exist today and some are sequenced work. Available now:

Web app

The full product — discovery, reading, timelines, the citation-first editor, the editorial workflow, the archive, oral histories and community.

In-workspace product surface

Composed into the wider workspace app shell and its administrative surfaces, so it sits alongside the other products a team already uses.

Floating AI assistant

Present on every screen and already aware of the story, topic, discussion, collection or archive item you have open.

Agent tools over MCP

Thirty-one domain tools plus twenty-one shared workspace modules, four read-only resources and two guided recipes, all acting under your own permissions.

Public read API

An unauthenticated GraphQL read surface over public-access archive items and oral histories, for partners, researchers and aggregators.

IIIF image endpoint

A Level-0 IIIF Image API with a compliant capability document, so heritage and museum viewers can render public material without an integration project.

Marketing website

Features, use cases, services, the AI page, security posture, FAQ, and the launch waitlist at theculturepost.com.

Documentation site

Product guides and API reference, with the contributor and editor handbooks still being written.

10 Front doors on the roadmap

Mobile apps

Reading, following, field capture of oral histories and approvals on the go, for iOS and Android.

Command-line tool

Terminal-native automation, bulk import, citation validation and scripting against the content API.

Node.js and Python SDKs

Typed programmatic access for integrations, and an idiomatic async Python client for research and digital-humanities pipelines.

Browser and editor extensions

Clip a primary source and attach a citation from any tab, and reach the citation-first editor from the editor you already write in.



11 Plans

Five tiers, priced on output rather than logins. Every user-facing capability is available from Free onward — tiers differ by capacity, governance, support and SLA. These are planning prices being finalised ahead of the 2026 launch.

FREE

\$0 / ₹0 forever

Independent writers, educators, small heritage groups. No credit card.

10 published stories, 5 topics, 2 contributor seats

2 collections, 50 archive items, 5 oral histories

1 GB storage, 500 API requests a day

Community support plus one ticket a month

STARTER

\$39/mo · \$31/mo yearly

Independent publications, faith communities and local history societies.

100 published stories, 25 topics, 5 contributor seats

15 collections, 1,000 archive items, 100 oral histories

25 GB storage, 10,000 API requests a day

Email support within 48 hours

GROWTH

\$149/mo · \$119/mo yearly

Growing cultural publications, museums and university departments.

1,000 published stories, 150 topics, 15 contributor seats

100 collections, 25,000 archive items, 2,500 oral histories

100 GB storage, 100,000 API requests a day

Priority support within 24 hours, 99.5% SLA

BUSINESS

\$499/mo · \$399/mo yearly

Established media houses, archives, religious organisations and cultural institutions.

10,000 published stories, 1,000 topics, 50 contributor seats

1,000 collections, 250,000 archive items, 25,000 oral histories

1 TB storage, 1,000,000 API requests a day

Priority support within 8 hours, 99.9% SLA

ENTERPRISE

From \$8/contributor/mo

Universities, national archives, broadcasters and multi-entity networks. Annual contract, invoiced.

\$6/contributor/mo annual floor at 50 or more contributors

Custom or unlimited quotas, enforced SSO with SAML or OIDC, SCIM

Dedicated tenancy, self-hosting and data-residency options

Custom SLA up to 99.95% and above, named contacts

11 How the pricing works

Four permanent regional bands

Pricing is set per band rather than converted from a single list. In India, Starter is ₹1,599/mo, Growth ₹5,999/mo and Business ₹19,999/mo, with GST charged separately. Lower bands also cover Eastern Europe and the Gulf, Latin America and South-East Asia, and much of South Asia and Africa.

Capacity add-ons rather than tier jumps

Packs extend a single dimension without moving you up a tier: story packs from \$19, topic packs from \$15, archive-item packs from \$29, oral-history packs from \$39, collection packs from \$29, citation packs from \$29, storage from \$15, API request packs from \$19, and contributor seats from \$49 for five.

No surprise overages

Quotas warn at eighty per cent and block new usage at one hundred. Nothing auto-purchases a pack on your behalf, and no usage is billed after the fact. Capability add-ons for analytics, audit retention and branding are priced separately and listed openly.

Services and infrastructure shapes

Implementation runs from a one-day quick start at \$2,500 to a two-to-four week standard engagement at \$12,000, with enterprise migrations quoted per project. AI allowances appear in every paid tier in bring-your-own-key terms — but bring-your-own-key is not supported today, so no AI spend figure is quoted.

12 Why teams choose it — evidence and rights

Sourcing is enforced at the claim, not offered as a plugin

A citation is a first-class record with a typed source kind, provenance and an editorial verification status, attached to the claim it supports. Most publishing tools treat a footnote as formatting; here it is the structure the product is built on.

Four rights tiers on every heritage item

Public, community, restricted and embargoed sit on each archive item and oral history, with restricted material reachable only through a separate elevated-access path. Sacred and community-only material has somewhere to live that is neither the open web nor a locked drive.

Verification status is a field, not a badge

Corroborated, disputed, retracted or unverified is assigned during editorial review and travels with the claim to the reader, so the strength of evidence is visible on the page rather than implied by the venue.

An editorial lifecycle with a sensitivity gate

Idea, draft, in review, published, archived — with only valid transitions permitted and a fact-and-sensitivity pass standing between a draft and publication, across six story formats.

12 Why teams choose it — dialogue, AI and governance

Interfaith dialogue is modelled, not moderated afterwards

A discussion can be marked an interfaith space, framed with a pinned prompt, carry perspective tags and four moderation states, and be locked when a thread needs to cool.

Open by default at the edges

A public unauthenticated read API and a Level-0 IIIF image endpoint mean heritage viewers can already render public material, and a researcher can query it, without an integration project.

AI that refuses instead of improvising

Answers are built only from your own attached sources, return the exact list they were given, declare their own uncertainty, and are refused outright when there is nothing to ground on.

Institutional governance from the first day

Single sign-on, a seeded role catalogue across eight resource types, workspace isolation, an audit timeline, quotas and billing arrive with the space.

Sealed material never leaves the archive

The restriction is a floor inside the AI layer rather than a permission check, so someone who can personally open a sealed testimony still cannot route it to a provider.

One data model, not two systems

A story and the manuscript it cites live in the same product, so publishing about a collection and stewarding it are not separate platforms to reconcile.

12 Honest answers to hard questions

You are pre-launch. Why would we commit now?

The editorial, archive, community and permission layers are built; the AI generation, review gate and run ledger are built without a screen in front of them; the reader-facing AI layer, mobile apps, SDKs and classroom tooling are not built. If your evaluation depends on something in the third group, wait.

The AI features you describe most confidently have no user interface.

Correct, and we would rather say so than show a mockup. The four grounded tasks, the review gate, the run ledger and the workspace-scoped permissions run in the service today and are reachable through the API and connected agents. The one AI surface you can use in the app right now is the floating assistant.

Our contributors could publish without review.

That is a real gap. Advancing a story through its lifecycle is currently enforced as the same permission as editing its body. A separate publish action is the fix and it is not shipped. In the meantime the mitigation is procedural: hold contributor accounts to review discipline, and use the recorded lifecycle history.

12 Collections, sacred material, procurement

We already run a digital-collections platform. Why add another system?

If your need is purely to catalogue and expose a collection, a mature collections platform will serve you well. TheCulturePost is for institutions that also want to publish about the collection — sourced narrative, editorial review, moderated public conversation — on the same data model as the material itself.

Our holdings include sacred and community-restricted material. How is that safe here?

Four access levels sit on each item and testimony; restricted holdings are reached through a separate elevated-access path; transcripts are encrypted at rest; the broad membership role is withheld from oral histories entirely. What is not built: community consent, takedown and access-request flows as first-class auditable workflows.

You have no SOC 2 report, so we cannot get this through procurement.

That is a fair blocker and we will not dress it up. There is no certification today. Practices are modelled on ISO 27001 and designed against SOC 2 criteria, and a data processing agreement is available now. For institutions whose policy cannot accommodate shared tenancy, dedicated, self-hosted and air-gapped options exist.

12 Languages, model keys, and trusting the AI

Our readers and contributors work in six languages and three scripts.

Partly served. Stories, archive items and oral histories record the languages they cover, and the interface is built for multiple languages with right-to-left support and locale-aware formatting. What does not exist is automatic translation or transliteration — that layer is designed and sequenced, not built.

We need to bring our own model key, or we cannot use the AI at all.

Not supported today. Provider keys are held on the vendor side across five model choices, with a per-request override, and a missing key returns a clear provider-unavailable error. The AI layer can also be switched off independently, so a workspace can adopt the platform and decline the AI entirely.

How do we know the AI is not quietly making things up anyway?

You check the record rather than trusting the claim. The list handed to the model is assembled from real records before the call and stored with the run. Every answer carries a self-declared uncertainty rating, and nothing publishes itself. What we do not have is an evaluation set, so no accuracy figure is claimed.

12 Where it sits against the alternatives

VS HERITAGE AGGREGATORS

Aggregators are unmatched at scale and at federating what institutions already hold. What they do not give you is a place to write about the material with claim-level sourcing, run editorial review, and host moderated conversation on the same data model.

VS ORAL-HISTORY PROGRAMMES

Established programmes have decades of practice and recording infrastructure we do not have — we do not even capture audio in-app. What we add is what happens after capture: a described, rights-tiered, citable record referenced inside a sourced story.

VS OPEN-SOURCE COLLECTIONS PLATFORMS

They are mature, extensible and free, and for a purely cataloguing brief they are a rational choice. Our argument is narrower: lifecycle, verification, moderated dialogue and a governed AI layer arrive as one product with roles, audit and billing in place.

VS LONG-FORM CULTURE PUBLICATIONS

They have editorial reputation, commissioning budgets and a readership. But they are publications, not platforms — a reader cannot query their archive, and a claim in an essay is not a structured record with a verification status attached.

VS ACADEMIC EXPLAINER NETWORKS

Their sourcing credibility is genuinely high, because the authors are the researchers. They do not offer heritage preservation, rights tiers on sensitive material, oral-history records, or a space an institution can run under its own brand.

VS GENERAL PUBLISHING TOOLING

It is more polished and far more widely supported, and its editor is nicer to use. It has no concept of provenance, of a rights posture on sacred material, or of an AI layer that refuses to answer without sources.

START TODAY

Culture, religion, values and history — thoughtfully explored

TheCulturePost launches in 2026 and is in waitlist mode today. If you steward a collection, run an editorial desk, or teach from primary sources, the useful conversation is the one that happens before launch.

- Join the waitlist to be first in line for the 2026 launch
- Read the features and the AI page, including the parts that are not built yet
- Institutions: talk to us about branded spaces, managed collections, dedicated tenancy or self-hosting
- Security and procurement questions go straight to the security address below

WEBSITE AND WAITLIST

theculturepost.com

WHAT IS BUILT

theculturepost.com/features

HOW THE AI WORKS

theculturepost.com/ai

INSTITUTIONS AND GENERAL ENQUIRIES

contact@theculturepost.com

RESPONSIBLE DISCLOSURE

security@theculturepost.com



TheCulturePost

TheCulturePost is pre-launch. Capabilities marked roadmap are not available today, and no adoption, revenue or accuracy claim is made anywhere in this document.