

UVA Archival AI Protocol Implementation Toolkit

Version 1.1, January 27, 2026

Forms, workflows, contract language, and operational resources

Created by Leo S. Lo

University of Virginia Library

**Developed with the UVA Special Collections and
Preservation staff**

2026 • CC BY 4.0

**This Implementation Toolkit is bundled with the Adoption Guide.
Both files share a single DOI.**



How to use this toolkit

This document contains the operational resources needed to implement the UVA Archival AI Protocol. It is designed for the people who do the work, archivists in any setting, librarians, digital services staff, and the administrators who support them.

The companion Adoption Guide explains the core rule, the retrieval-versus-training distinction, the three adoption levels, and the public statement template. It is recommended that the Adoption Guide be read first.

The Adoption Guide and this Toolkit are two parts of a whole. The Guide covers sections 1-5; this toolkit covers sections 6-16.

What's in this toolkit:

Section 6: Default restrictions · Section 7: Review workflow · Section 8: Conditions library · Section 9: Intake form · Section 10: Decision log · Section 11: Provenance standard · Section 12: Model clauses · Section 13: Implementation checklist · Section 14: Attribution · Section 15: Staff FAQ · Section 16: Contact

6. Default restrictions

Use these as defaults. Institutions can tighten or expand based on local policy and law. The key organizing principle is the distinction between retrieval-based and training-based AI uses (refer to the Adoption Guide for these distinctions).

6.1 Uses not permitted by default

- Training, fine-tuning, or embedding at scale using archival materials, unless explicit review confirms item-level provenance, meaningful attribution, and enforceable institutional control including the right to stop further use.
- Uploading archival materials to third-party tools where the institution cannot control retention, reuse, or downstream training.
- Any use that conflicts with donor restrictions, deed of gift and transfer terms, community protocols, or privacy commitments.
- Any use that risks revealing personal data, sensitive community information, or materials with re-identification risk.
- Any use that enables surveillance, profiling, targeted harassment, or harmful practices toward individuals or communities represented in the archival material.

6.2 Uses that may be permitted with conditions

Retrieval-based uses: Transcription, translation, description assistance, summarization, or discovery-oriented analysis where source items remain under institutional control, outputs preserve attribution, and sensitive content is not exposed.

Controlled internal models: On-premises or institution-hosted models for metadata creation, search, collection analysis, preservation assessment, or accessibility features, where the organization can delete data and shut down the system.

Scoped research use: Research using retrieval or time-limited analysis under institutional oversight, with a data management plan submitted and approved.

7. Review workflow and decision path

This workflow branches on the fundamental distinction the Protocol draws between retrieval-based and training-based AI. It is designed to be repeatable and auditable.

7.1 Workflow steps

1. **Intake.** Requester submits the intake form (Section 9) with a description of materials, purpose, tool or vendor, intended outputs, and whether the proposed use is retrieval-based or training-based.
2. **Classification.** Reviewer confirms whether the request involves retrieval/analysis or training/fine-tuning/adaptation. If the requester's self-classification is unclear or incorrect, the reviewer reclassifies. This step determines which decision path applies.

Classification quick test. If the reviewer is unsure how to classify a request, ask these five questions. If the answer to any is yes or unknown, classify as training-based and route to Path B.

1. Does the tool send our materials to an external server the institution does not control?
2. Do the vendor's terms of service include language about training on, learning from, or improving models using uploaded content?
3. Will our materials become part of a model that exists independently of our infrastructure — one we cannot delete or shut down?
4. After the project ends, would it be technically impossible to remove our materials from the system?
5. Could the requester or vendor redistribute a model or dataset that contains our materials to third parties?

3. **Rights and ethics check.** Review rights status, donor restrictions, community obligations, sensitivity flags, and whether the request is reversible. If rights or ethics are unclear, decline or seek clarification before proceeding.

Indigenous and community-held materials. For collections involving Indigenous knowledge, Native American materials, or materials from marginalized communities, the rights and ethics check must go beyond donor consent and institutional control. The right of communities to determine how their knowledge is used is a distinct obligation. Reviewers should consult the Protocols for Native American Archival Materials and the CARE Principles for Indigenous Data Governance (Collective Benefit, Authority to Control, Responsibility, Ethics) as additional review frameworks. Where community consultation is warranted, AI use should not proceed until that consultation is complete.

4. **Decision path.** For retrieval-based requests, apply Path A. For training-based requests, apply Path B. Assign a decision: allow, allow with conditions, or do not allow.
5. **Conditions.** If conditions are required, issue an approved conditions letter using the conditions library (Section 8), including retention limits, non-training terms, and the institution's right to stop further use.
6. **Record and schedule.** Record the decision in the decision log (Section 10). Store supporting documents. Schedule a follow-up review date for time-bound approvals.

7.2 Path A: retrieval and controlled internal models

Confirm all of the following before approving:

- The archival organization, or a trusted partner under contract, can remove its data and shut down the system if needed.
- Outputs maintain clear links back to original items or collections, consistent with the provenance standard in Section 11.
- The project advances discovery, preservation, research, teaching, or community benefit.
- Sensitive or restricted materials are handled according to existing access policies.

If all are confirmed, approve in principle and apply appropriate conditions. If any cannot be confirmed, revise the project scope or decline.

7.3 Path B: training or adapting models

This path applies to any use that will absorb knowledge from archival materials into model weights. Apply the core rule directly:

- **Provenance:** Can the partner maintain item-level provenance and meaningful attribution in practice?
- **Control:** Can the archival organization retain a contractually enforceable control to stop further use, including the right to demand cessation of ingest and, for narrow models, decommissioning?

If the answer to either is no, the default decision is to decline. If both appear possible, classify the project:

Narrow model (trained only or mainly on defined archival collections): Require strong decommission rights, detailed documentation, provenance to the standard in Section 11, and appropriate compensation or benefit sharing.

Mixed-corpus or foundation model (an archival collection may be one source among many with multiple items within a collection): Recognize that full decommissioning is unlikely. At minimum require: no new training runs using the organization’s materials after notice, removal from retrieval and indexes on request, clear documentation, and benefit sharing. If these cannot be met, do not proceed.

When the answer is no. When a request is declined, communicate the decision clearly and leave the door open for a revised proposal. Suggested language:

“Thank you for your interest in our collections. After review under the UVA Archival AI Protocol, we are not able to support this use as proposed because [the core rule cannot be met / item-level provenance cannot be demonstrated for the requested materials / the institution cannot retain enforceable control over further use / the proposed use conflicts with donor restrictions or community obligations]. We welcome a revised proposal that addresses these concerns. If you would like to discuss alternative approaches, including retrieval-based methods that may be supportable, please contact [name].”

7.4 Final questions before any approval

Before agreeing to any AI use, confirm:

1. The use respects donors, creators, records custodians, communities, and ethical obligations.
2. The archival organization keeps real control — not just symbolic promises.

3. The project can be explained clearly to staff, donors, records creators, and the public.
4. Benefits to the organization and its communities justify any risks.
5. The decision and conditions are recorded in the decision log.

If any answer is no, tighten conditions or decline.

8. Conditions library

Select conditions as needed. Use plain language. Add local legal review where required.

8.1 Core conditions

No training. Materials and derived outputs may not be used for model training, fine-tuning, or improving any AI system.

No retention. Materials must be deleted by a specified date, including from backups, caches, and logs where feasible.

No redistribution. Materials may not be shared outside the approved project team.

Attribution and provenance. Any permitted output must include provenance information consistent with the standard in Section 11, including collection citation and restriction statements.

No sensitive exposure. Outputs may not reproduce sensitive personal data or restricted content.

Audit and attestation. Requester must provide an attestation of compliance and allow reasonable verification.

Right to stop. The institution may require cessation of all further use of archival materials upon written notice, including cessation of ingest into new training runs, removal from retrieval indexes, and deletion of derived datasets.

8.2 Optional conditions

- On-premises processing only, using institution-controlled infrastructure.
- Limited corpus scope: specific folders or items only, with item-level tracking.
- Human review of all outputs before any public release.
- Time-bound approval with re-review required for extensions.
- Security controls for data transfer and storage, including encryption at rest and in transit.

- Decommissioning rights: For narrow models, the institution may require decommissioning or destruction of the model if permission is withdrawn.
- Compensation or benefit sharing: fees, digitization support, infrastructure contributions, tool access, or other structured benefits.

9. Intake form template

Deploy this form through your institution’s preferred method (web form, shared document, ticketing system). The form forces requesters to classify their proposed use as retrieval-based or training-based at the outset, which determines the review path.

Adapt field labels and instructions to your local context.

AI use request: Intake form

Fields marked with an asterisk () are required.*

Requester information

Name* _____

Requesting entity* _____

Email* _____

Date of request* _____

Project description

Project title*

Purpose and goals*

Describe what you intend to accomplish and why.

Archival materials requested*

Identify specific collections, series, items, or item types.

Volume of materials*

Approximate number of items, pages, or files.

AI use classification

Type of AI use*

Select one:

- Retrieval-based (search, analysis, transcription, summarization — source items remain under institutional control)
- Training-based (model training, fine-tuning, adaptation, embedding into model weights)
- Unsure (reviewer will classify)

AI tool or vendor*

Name the specific tool, platform, model, or vendor.

Where will processing occur?*

Select one:

- On-premises / institution-controlled infrastructure
- Third-party cloud service (specify provider)
- Vendor-hosted platform (specify)

Will materials or outputs leave institutional control?*

- No
- Yes (explain)

Intended outputs and use

Describe intended outputs*

What will be produced? Transcripts, metadata, summaries, a trained model, embeddings, etc.

Will outputs be public?*

- No, internal only
- Yes (describe distribution)

Retention period*

How long do you need access to the materials and outputs?

Rights, restrictions, and sensitivity**Are you aware of any donor restrictions, privacy concerns, or community obligations?***

- No
- Yes (describe)
- Unsure

Does the material contain personally identifiable information or sensitive content?*

- No
- Yes (describe)
- Unsure

Acknowledgment**Requester acknowledgment ***

I understand that this request will be reviewed under the UVA Archival AI Protocol. I agree to comply with any conditions issued as part of the approval. I understand that approval is not guaranteed.

Signature: _____

Date: _____

10. Decision log template

Maintain this log as a shared spreadsheet or database. Each row records one review decision. The log serves as the institutional record of how AI requests were handled and supports audit, learning, and governance review.

Getting started at low volume. The template below shows all available fields. Adoption Level 1 requires no decision log. At adoption Level 2, the minimum viable log is a shared spreadsheet with five fields: request ID, date, requester, AI use type (retrieval or training), and decision. The remaining fields become required at adoption Level 3 or when volume justifies them. Even a

single row in a spreadsheet counts as a functioning decision log. Do not let the template's comprehensiveness become a barrier to starting.

Minimum fields are shown below. Fields marked with an asterisk (*) are required at adoption Level 2. Add columns for local needs (e.g., collection sensitivity rating, legal review date).

Field	Description	Example entry
Request ID*	Sequential identifier.	<i>AAIP-2026-001</i>
Date received*	Date the intake form was submitted.	<i>2026-02-15</i>
Requester*	Name, title, and affiliation.	<i>Dr. J. Smith, History Dept.</i>
Project title	Short title from the intake form.	<i>Civil War Letters Transcription</i>
Collections in scope	Collections, series, or items requested.	<i>MSS 1234, Series II, Boxes 3-7</i>
AI use type*	Retrieval-based or Training-based.	<i>Retrieval-based</i>
Tool or vendor	Name of AI system or provider.	<i>Institution-hosted Whisper model</i>
Decision path	Path A (retrieval) or Path B (training).	<i>Path A</i>
Decision*	Allow / Allow with conditions / Do not allow.	<i>Allow with conditions</i>
Conditions issued	List conditions applied from the library.	<i>No retention beyond 2026-08-15; attribution required; no redistribution</i>
Reviewer(s)	Name(s) and role(s) of reviewer(s).	<i>A. Lee, Head of Special Collections</i>
Decision date	Date the decision was made.	<i>2026-02-22</i>
Follow-up date	Scheduled re-review for time-bound approvals.	<i>2026-08-01</i>
Notes	Any additional context or rationale.	<i>Donor restrictions reviewed; no conflicts.</i>

11. Provenance standard for AI outputs

This section provides a practical, tiered version of the Protocol’s Appendix B (Minimum Provenance Standard for Citations and Logs). The full Appendix B should be consulted for detailed implementation, especially for public-facing or vendor-hosted uses.

11.1 Core provenance rule

If the AI system cannot cite the specific archival source that supports an output, it should not produce that output from archival materials—or should produce only the portions it can cite.

11.2 Tiered provenance approach

Most archival organizations hold materials at varying levels of description and digitization. The provenance standard is tiered to reflect this reality. Higher tiers are required as uses move from internal to external and from retrieval-based to training-based.

Tier	Material description level	Provenance standard	Permitted AI uses
Tier 1 Full	Fully described, digitized materials with persistent identifiers and item-level metadata.	Full provenance standard: all required citation fields (collection ID, item ID, location within item, holding organization, access link).	All permitted uses under the Protocol, including public-facing outputs and external research collaborations.
Tier 2 Container	Box-and-folder or series-level description. No item-level metadata or persistent identifiers.	Container-level citation: collection ID, series, box, folder, and holding organization. Flag that item-level provenance is not yet established.	Internal use and approved research use. Public-facing outputs and external use require item-level provenance to be established first.
Tier 3 Minimal	Unprocessed or minimally described materials. No finding aid or only accession-level record.	Collection-level citation only: collection ID, holding organization, and a note that the material is minimally described.	Internal discovery and description workflows only. Outputs are treated as draft, labeled as AI-generated, and subject to archival review before any release.

Note: Tier 3 turns the provenance gap into a feature. AI-assisted description of minimally processed collections is one of the clearest mission-aligned use cases under the Protocol, as long as outputs are flagged as AI-assisted and reviewed by an archivist before publication. Where item-level provenance cannot yet be maintained, AI use is permitted for internal workflows that improve description.

11.3 Required citation fields (provenance tier 1)

Every AI output that asserts facts, quotes, summarizes, or describes archival content at Tier 1 must include:

- **Collection ID and title.** Identifies the archival collection.
- **Item ID.** Persistent identifier preferred. Cite at the item level whenever the claim depends on a specific item.
- **Location within item.** Page number, timestamp, frame, or segment ID, when available.
- **Holding organization name.** The archival organization that holds the material.
- **Access link.** Public URL when available; internal resolver link for restricted materials.

11.4 Key rules

- Cite at the item level for specific claims. Use collection-level citation only for purely contextual statements.
- For multi-item synthesis, provide multiple item citations, not a single broad citation.
- Do not cite an item unless it directly supports the statement.
- For direct quotes and transcriptions, cite the exact source item and location within the item.
- If transcription is uncertain, label the uncertainty and avoid definitive claims.
- If the system cannot produce the required fields for the applicable tier, it must either decline to answer from archival sources or provide only the portions it can cite, with a clear note that other parts were withheld.

11.5 Minimum logging

For each AI interaction involving archival materials, capture when feasible: date and time, system name and version, user category (staff, researcher, public), query text (or hashed query for sensitive content), retrieved item and collection IDs, citations shown to

the user, and output text where retention is permitted. Limit log access to staff with a clear need.

12. Model clauses

These clauses are starting points. Legal counsel should review and adapt to local context before use. The Protocol's Part VIII contains additional clause variants; adopting institutions should consult both sources.

12.1 Deed of gift clause (standard)

The Donor/Creator grants to the Archives the right to use the archival materials, and any digital reproductions of them, in AI-supported tools and services that the archival organization controls, including search, description, preservation, and accessibility functions.

The archival organization will not authorize the use of this collection to train or adapt general-purpose AI models that cannot maintain item-level provenance and meaningful attribution, or that prevent the archival organization from exercising contractually enforceable control over continued use, including the right to stop further use.

Any broader AI training use of this material will require a separate written agreement that specifies scope, conditions, and protections for the Donor/Creator and the archival organization.

12.2 Deed of gift clause (restrictive)

The archival materials may not be used in any AI training, model development, or related computational process without the prior written consent of the Donor, Creator, or their representative. In the case of an institutional transfer of records, prior written consent of the administrative body is required.

The archival organization will not authorize the use of these materials in training general-purpose AI models under any circumstances.

Any request for AI-related use must specify whether it is retrieval-based or training-based and must be approved in writing by both the Donor/Creator/Records Custodian (or representative) and the archival organization.

12.3 Bridge guidance for pre-existing agreements

Most archival organizations hold collections acquired under deeds of gift, transfer agreements, and purchase terms that predate generative AI and say nothing about AI use. This section provides guidance for applying the Protocol to those collections without retroactively amending existing agreements.

Default position. For collections acquired under agreements that predate the Protocol, the organization applies the default restrictions in Section 6 as institutional policy. The

core rule functions as a supplemental policy, governing uses that the original gift agreement never contemplated. No retroactive amendment to existing deeds is required for the default position to take effect, because the default position is restrictive — it declines to authorize training, it does not grant new rights that require donor consent.

Broad legacy language. Where existing deeds contain broad language that could be read to permit AI training — for example, language granting the institution unrestricted use for “any educational purpose” or “all forms of reproduction” — the organization should flag those agreements for legal review and apply the default restrictions pending that review. Broad historical language was not drafted with irreversible AI training in mind and should not be treated as blanket authorization.

Review during workflow. For AI use requests involving collections under legacy agreements, the reviewer should check the deed language as part of the rights and ethics step (Step 3 in the workflow) and note any ambiguity in the decision log. Where the deed is silent on AI, the Protocol’s defaults govern. Where the deed language is genuinely ambiguous, consult legal counsel before proceeding.

Proactive outreach. For high-value or high-risk collections, the organization may choose to contact living donors or their representatives to discuss AI use preferences and, where appropriate, to execute an addendum using the standard or restrictive clause. This is recommended but not required for the default restrictions to apply.

12.4 Vendor clause

Vendor agrees that any content provided by the archival organization under this agreement may be used solely to deliver the contracted services and may not be used for generative AI training, fine-tuning, or model improvement. Vendor may not retain, reuse, or disclose the content or derived outputs beyond the contract term.

Any approved AI use of archival collections must either be retrieval-based or involve narrow models that the organization can cause to be decommissioned if permission is withdrawn. Vendor must implement appropriate security safeguards, maintain documentation sufficient to verify scope of use, and provide written attestation of deletion upon request.

The archival organization retains the right to require Vendor to stop any further use of archival material in AI training, evaluation, or retrieval upon written notice.

12.5 Researcher access agreement clause

Researcher agrees to use approved archival materials only for the stated project purpose and to comply with all conditions issued under the UVA Archival AI Protocol. Researcher will not upload non-public materials to third-party tools without explicit approval. Researcher will not use the materials or derived outputs for model training, fine-tuning, or dissemination beyond the approved scope. Any training or adaptation of models that will leave the organization's control or be widely redistributed is subject to separate review and agreement.

Institutions should incorporate these terms into researcher registration forms, reading room policies, and any digital access portals so that agreement is captured at the point of access.

13. Implementation checklist

13.1 Adoption level 1: under 30 days

- Choose a canonical URL for the Protocol and host the PDF.
- Publish the public statement (see Adoption Guide) with the core rule language and required attribution line.
- Adopt the default restrictions (Section 6) and name the request contact point.
- Deploy the intake form (Section 9) through your preferred channel (web form, shared document, or email template).
- Brief your immediate team on the core rule and where to route AI requests.

13.2 Adoption level 2: under 60 days

- Confirm the review workflow (Section 7) and assign at least two reviewers by role.
- Connect the intake form to a shared inbox or ticketing system.
- Start the decision log (Section 10) using the template. At minimum, track the five required fields: request ID, date, requester, AI use type, and decision.
- Create a conditions letter template using the conditions library (Section 8).
- Process at least one request through the full workflow to test the system.
- Add AI clauses (Section 12) to new deeds of gift, institutional transfers, and vendor agreements.

13.3 Adoption level 3: under 120 days

- Create a standing review group charter and set meeting cadence (quarterly recommended).
- Implement light audit and attestation: at least annual sampling of approved uses.
- Set an annual review date for the Protocol adoption and publish version updates with a change log.
- Conduct a short internal training session for all staff who handle archival AI requests. Use the Staff FAQ (Section 15) as a training resource.
- Review high-risk or high-value collections and existing agreements for AI implications. Apply the bridge guidance in Section 12.3 to flag legacy deeds with ambiguous language.
- Add a notice to institutional transfer documentation.
- Audit existing vendor contracts for AI training clauses, especially platforms that hold digitized collections or manage digital assets. Prioritize digital asset management systems, integrated library systems, digitization service providers, and cloud storage or hosting platforms. For each vendor, determine whether current terms include AI training rights and flag any that do for renegotiation or addendum using the vendor clause in Section 12.4.

14. Attribution and reuse

Required attribution line

All adopters must include this attribution in their public statement and any published adaptation:

This work is based on the UVA Archival AI Protocol, created by Leo S. Lo and developed with the UVA Special Collections and Preservation staff, University of Virginia Library. Canonical URL: <https://doi.org/10.18130/5dqf-9w86>.

Suggested citation

Lo, Leo S. The UVA Archival AI Protocol: A provenance, consent, and control standard for generative AI use of archival collections. Developed with the UVA Special Collections and Preservation staff. University of Virginia Library. Version [X.X], [Year]. Canonical URL: <https://doi.org/10.18130/5dqf-9w86>.

License and badge

Recommended license for broad adoption with required attribution: Creative Commons Attribution 4.0 International (CC BY 4.0).

Badge text: Adopter of the UVA Archival AI Protocol. Learn more at: <https://doi.org/10.18130/5dqf-9w86>.

15. Staff FAQ

These answers are designed for archival staff who receive questions in reading rooms, from researchers, or from colleagues. Each answer is grounded in the retrieval-versus-training framework. Use this section for staff training and as a quick-reference guide.

A researcher wants to use ChatGPT to help read a handwritten letter from our collection. Is that allowed?

It depends on whether the material is public or restricted. If the researcher is working with publicly available digitized material, we cannot prevent them from uploading it to third-party tools, but we can recommend practices that maintain attribution and note that commercial AI tools may train on uploaded content. If the material is non-public or restricted, our default position is no: Uploading non-public archival materials to third-party tools where we cannot control retention, reuse, or downstream training is not permitted without explicit review. Direct the researcher to the intake form and suggest they ask about institution-hosted transcription tools as an alternative.

Can we use AI transcription on our oral histories?

Yes, this is a core permitted use under the Protocol — as long as the tool is retrieval-based and we control the infrastructure. On-premises or institution-hosted transcription tools (for example, a locally hosted Whisper model) are straightforward Path A approvals. If the transcription service is cloud-based, check whether the vendor's terms include AI training rights on uploaded content. For oral histories with donor restrictions, cultural sensitivity, or community obligations, flag those for additional review during the rights and ethics check. Register the project in the decision log.

A vendor says their new search feature uses AI. Do we need to review that?

Yes. The key question is whether the feature is retrieval-based (searching our content at query time, where our data stays under our control) or training-based (using our content to improve or train the vendor's model). Ask the vendor directly and get the answer in writing. If they are training on our collections, this is a Path B scenario and the vendor

clause in Section 12.4 applies. Many platforms are quietly adding AI features to existing products — make this a standard question in vendor renewals and contract reviews.

A faculty member wants to build a custom GPT using our finding aids. How do we handle that?

Finding aids are intellectual content in scope under the Protocol. Start by classifying the use. If the faculty member is building a retrieval-based tool that searches finding aid text at query time (essentially a specialized search interface), this is likely Path A: Confirm they can maintain attribution, the data stays under institutional control or an approved platform, and sensitive content is handled appropriately. If they are fine-tuning a model on our finding aids so the model absorbs that knowledge into its weights, this is Path B and requires the full training-use review. In either case, the researcher should submit the intake form, and the project should be registered in the decision log.

Do we need to review every use of AI by our own staff?

At Level 1, the main requirement is that AI requests involving archival collections are routed through the contact point. At Level 2, internal projects using AI with archival materials should be registered in the decision log. The bar is intentionally low — a single row in the spreadsheet with the five required fields is enough. The point is institutional awareness and a light audit trail, not bureaucratic burden. Routine uses of general-purpose AI tools for everyday work tasks (drafting emails, summarizing meeting notes) that do not involve archival collections are outside the Protocol’s scope.

A donor asks whether their collection will be used to train AI. What do I say?

You can tell them clearly: “We will not let your material disappear into AI systems we cannot control. We follow the UVA Archival AI Protocol, which means we only support AI projects that keep provenance clear and that we can change our mind about later. We will not authorize the use of your collection to train general-purpose AI models.” If the donor wants to restrict AI use more tightly, offer the restrictive deed of gift clause (Section 12.2), which requires the donor’s written consent before any AI-related use. All new deeds of gift should include one of the AI clauses from Section 12.

We have thousands of existing deeds of gift that say nothing about AI. What do we do?

You do not need to retroactively amend existing deeds. The Protocol’s default restrictions (Section 6) function as institutional policy and apply to all collections, including those acquired under legacy agreements. Because the default position is restrictive — it blocks training, it does not authorize it — no new donor consent is required. For collections where legacy deed language is broad enough that it could be

read to permit AI training, flag those for legal review. For high-value or high-risk collections, consider proactive outreach to living donors. See Section 12.3 for detailed bridge guidance.

A company contacts us and wants to license our digitized collections for AI training. What do we do?

This is a Path B scenario and the highest stakes request type. Do not handle it at the departmental level alone. Route it to the institution's senior leader and legal counsel immediately. The core rule applies in full: The company must demonstrate item-level provenance and meaningful attribution in practice, and the institution must retain contractually enforceable control to stop further use. For broad commercial training of general-purpose models, these conditions are extremely difficult to meet. Apply the vendor clause (Section 12.4), the Right to Stop condition, and the compensation or benefit-sharing condition from Section 8. Any agreement requires a formal written contract reviewed by legal counsel. If the company cannot meet the core rule, decline. The Protocol's default position exists precisely for this scenario.