

Framework whitepaper

The full framework as a formal whitepaper — abstract, the seven documentation layers, artifacts, visibility policy, and operating model. Save it as a PDF to share or attach.

WHITE PAPER

A Campus AI Registry & Governance Framework

A layered model for documenting, governing, and disclosing artificial intelligence in higher education.

AUTHOR

Joe Sabado · CampusAIXchange.com

PUBLISHED

June 2026

VERSION

1.0

CLASSIFICATION

Public

LICENSE

CC BY 4.0

AUDIENCE

AI strategists & leadership

ABSTRACT

Colleges and universities now run many AI tools at once — institution-managed platforms, vendor suites, embedded enterprise AI, CRM and LMS features, and local research environments. This whitepaper presents a reusable framework that documents AI at every level: institutional governance, context-specific risk profiles, user-facing service cards, deployment-level system cards, model documentation, and the internal evidence behind governance decisions. A four-label visibility scheme determines what is published and what stays internal. The result is both a transparency tool for the community and a governance operating model for reviewers and auditors.

CONTENTS

1. Introduction
2. The documentation hierarchy
3. Core artifacts
4. How artifacts relate
5. Implementation types
6. NIST AI RMF profiles
7. Visibility & publication policy
8. Operating model
9. Adoption & peer context
0. Conclusion
1. Appendix — Key terms

1. Introduction

Most institutions did not adopt AI through a single decision. Tools arrived through procurement, vendor suites, research computing, and features switched on inside existing systems. The documentation that should accompany those tools — who owns them, what data they may handle, what risks they carry — tends to be scattered or missing entirely.

This framework gives that documentation a single, layered structure. It moves from broad institutional policy down to specific models, and from user-friendly guidance toward technical evidence. Each layer answers a narrower question than the one above it, and every artifact carries an explicit audience and visibility label so that publishing decisions are deliberate rather than accidental.

2. The documentation hierarchy

The framework is organized as seven layers, narrowing from institution-wide rules to the evidence behind a single deployment.

LAYER	WHAT IT CONTAINS
1 • Institutional governance	Strategy, AI principles, acceptable use, data classification, academic integrity, procurement and privacy/security/accessibility requirements, and the consolidated AI risk registry that rolls up risk across the portfolio.
2 • NIST AI RMF profiles	Context-specific overlays that adapt the NIST AI Risk Management Framework to particular settings — context profiles that govern use, and platform profiles that govern shared services.
3 • Service catalog & use-case registry	Two paired institutional indexes — a catalog of supported services (the supply side) and a registry of approved and proposed AI use cases (the demand side) — each with scope, audience, and allowed data classes.
4 • Service cards	User-facing documentation: what a tool is, for whom, and under what rules.
5 • System cards	Deployment-level documentation: architecture, integrations, controls, and governance boundaries.
6 • Model cards	Vendor model cards or internal summaries covering intended use, limits, and local context.
7 • Supporting evidence	Per-deployment risk registers, reviews, data-flow diagrams, evaluations, incident records, and due diligence — the local evidence the consolidated registry rolls up.

3. Core artifacts

The artifacts are complementary, not duplicative — each answers a different governance question. Three card types do most of the connective work:

- **Service cards** are the main user-facing documents, explaining scope, audience, approved use, and data guidance.
- **System cards** sit beneath or beside service cards when a service is a configured institutional deployment. For a multi-model platform, the system card also records the model catalog, routing logic, and change management — making model selection a documented governance decision rather than a configuration detail.
- **Model cards** sit below the system level and describe the actual models used, whether vendor-published or internally summarized.

Around these sit institutional principles, acceptable-use guidance, RMF profiles, the service catalog, and the supporting evidence layer — each with a defined owner, audience, and visibility.

4. How artifacts relate

A simple rule keeps the hierarchy coherent as it grows: move from broad policy to narrow implementation, and from user-friendly guidance to technical evidence.

Rule of thumb. If it is a tool users choose from, create a service card. If it is a configured institutional deployment, add a system card. If it is a specific model or a local model-selection decision, add a model card or internal model summary.

A shared institution-managed platform — one service offering several models that users and developers can choose from — is governed along two complementary lanes, continuously rather than once:

- **Platform governance** covers the shared service itself: ownership, approved data classes, the model catalog and routing, security controls, logging, and change management (mainly layers 1, 2, 5, 6, 7).
- **Tool & use governance** covers what people build on top and how the community may use it: downstream tools, use cases, eligible groups, and required approvals (mainly layers 2, 3, 4, 7).

Each lane has an institutional index at layer 3. The **service catalog** records the services on offer; a parallel **use-case registry** records the uses the community is approved for or has requested. A use case is mapped against the relevant context profile to decide approval, and if it hardens into a configured deployment it graduates into its own service or system card.

Two registries, two altitudes. Risk is documented twice over: a **per-deployment risk register** sits in the evidence layer beneath each system card, while a **consolidated AI risk registry** at the governance layer aggregates those entries into one portfolio view — the documented output of the RMF's Map, Measure, and Manage functions. The use-case registry is the demand-side companion to the service catalog, governed by the same context profiles.

5. Implementation types

The same framework documents very different categories of AI offerings, each tending to require a particular combination of artifacts.

TYPE	LIKELY ARTIFACTS
Institution-managed AI platforms	Service + system card + model references
Vendor AI assistants & productivity suites	Service card; optional tenant note
Embedded enterprise AI	Service card + often a system card
CRM, outreach & engagement AI	Service + system card when configured deeply
LMS, classroom & assessment AI	Service card; optional LMS note
AI-as-a-Service & API gateways	System card + model references + service summary
Local & research AI environments	System card + internal model summaries + risk evidence

Institution-managed platforms with several selectable models, APIs, and extensible tools warrant the most attention here: they require both platform governance and tool/use governance, kept under continuous review as models, modalities, and integrations change. A sandbox is one mode of such a platform — an experimentation space governed by the same system card — rather than a separate implementation type.

6. NIST AI RMF profiles

The middle of the stack is anchored by profiles built on the NIST AI Risk Management Framework. The framework's four functions — **Govern, Map, Measure, Manage** — are voluntary and not specific to higher education, but they can be tailored into profiles for particular contexts: teaching and learning, research, administration, student engagement, procurement, and agentic or cloud AI.

Profiles turn a general risk framework into context-specific guidance and govern how services in each setting are reviewed and approved. Full profiles are typically internal; a public summary communicates the expectations that apply to each setting.

Layer 2 carries two kinds of profile, matching the two governance lanes: **context profiles** (teaching and learning, research, administration, student engagement) that govern how AI is *used*, and **platform profiles** for institution-managed platforms, cloud AI, or agentic systems that govern the shared *service* itself.

7. Visibility & publication policy

Not every document should be public. Each artifact carries one of four visibility labels, so transparency means disclosing the right thing to the right audience — not publishing everything.

LABEL	MEANING
Public	Safe for open publication.
Institution-wide	Visible to all credentialed campus users.
Restricted Internal	Limited to approved internal stakeholders.
Confidential	Limited due to sensitive security, privacy, legal, or incident content.

Sensible defaults follow: publish user guidance, keep control evidence internal, treat model documentation as mixed-visibility, and publish summaries when full documents are too sensitive.

8. Operating model

A lightweight operating model keeps the framework maintainable. Every artifact opens with a short metadata header — title, document type, primary audience, visibility, owner, approver, version, and review date — and moves through a five-step review:

1. Author proposes the document type and visibility.
2. Service owner confirms scope and intended audience.
3. Privacy, security, and governance reviewers validate the visibility choice.
4. Approver confirms publication location and access controls.
5. Visibility is rechecked whenever the service changes materially.

What counts as a material change. For a shared platform, review is continuous rather than one-time: enabling a new model, modality, tool, or integration — or opening a new category of downstream use — is a material change that re-triggers the workflow before it goes live.

9. Adoption & peer context

The framework is not speculative. Peers commonly publish AI principles, acceptable-use and academic-integrity guidance, tool catalogs, and vendor model-card links. Less common — and where this framework adds the most — are formal RMF profiles by domain, reusable service-card templates, system cards for institution-run deployments, and unified visibility rules across all artifacts. The framework assembles existing practice into a more complete institutional model.

10. Conclusion

Adopted together, these layers let the campus community understand what AI exists and how to use it, while giving reviewers and auditors enough internal evidence to assess risk, controls, and accountability. It is both a transparency framework and a governance operating model — a registry of the institution's AI, documented once and disclosed appropriately.

A. Key terms

A short, neutral vocabulary for documenting institution-managed platforms.

TERM	MEANING
Institution-managed AI platform	A shared institutional service providing one interface to multiple AI models, often with multimodal capabilities and API-enabled extensibility.
Multi-model	The platform offers several underlying models — and sometimes providers — that users and developers can choose from through the shared service.
Model catalog & routing	The enabled models and providers, what each is approved for, and how the platform decides which model handles a request.
Platform governance	Governance of the shared service itself: ownership, controls, the model catalog, logging, and change management.
Tool & use governance	Governance of what people build on the platform and how the community may use it — downstream tools, use cases, eligible groups, and approvals.
Continuous governance	Ongoing review of a production platform, rechecked on every material change rather than approved once.
Material change	A new model, modality, tool, or integration, or a new category of downstream use — each re-triggers review before going live.
Governance boundary	The data classes and use cases a service is approved for — and what is out of scope and where it is routed instead.
Use-case registry	The demand-side index of AI uses the community is approved for or has proposed — the companion to the service catalog, governed by context profiles.
AI risk registry	The consolidated, portfolio-level roll-up of AI risks, aggregating the per-deployment risk registers held as supporting evidence.

A Campus AI Registry & Governance Framework — by Joe Sabado (CampusAIXchange.com), June 2026, version 1.0.

Licensed under **CC BY 4.0**. You may share and adapt this work, including at your own institution, with attribution to Joe Sabado.

This is a personal project. It is not associated with, supported by, or endorsed by the author's employer; views expressed are the author's own.

Use **Download (Print / Save as PDF)** and choose “Save as PDF” as the destination to download the whitepaper.