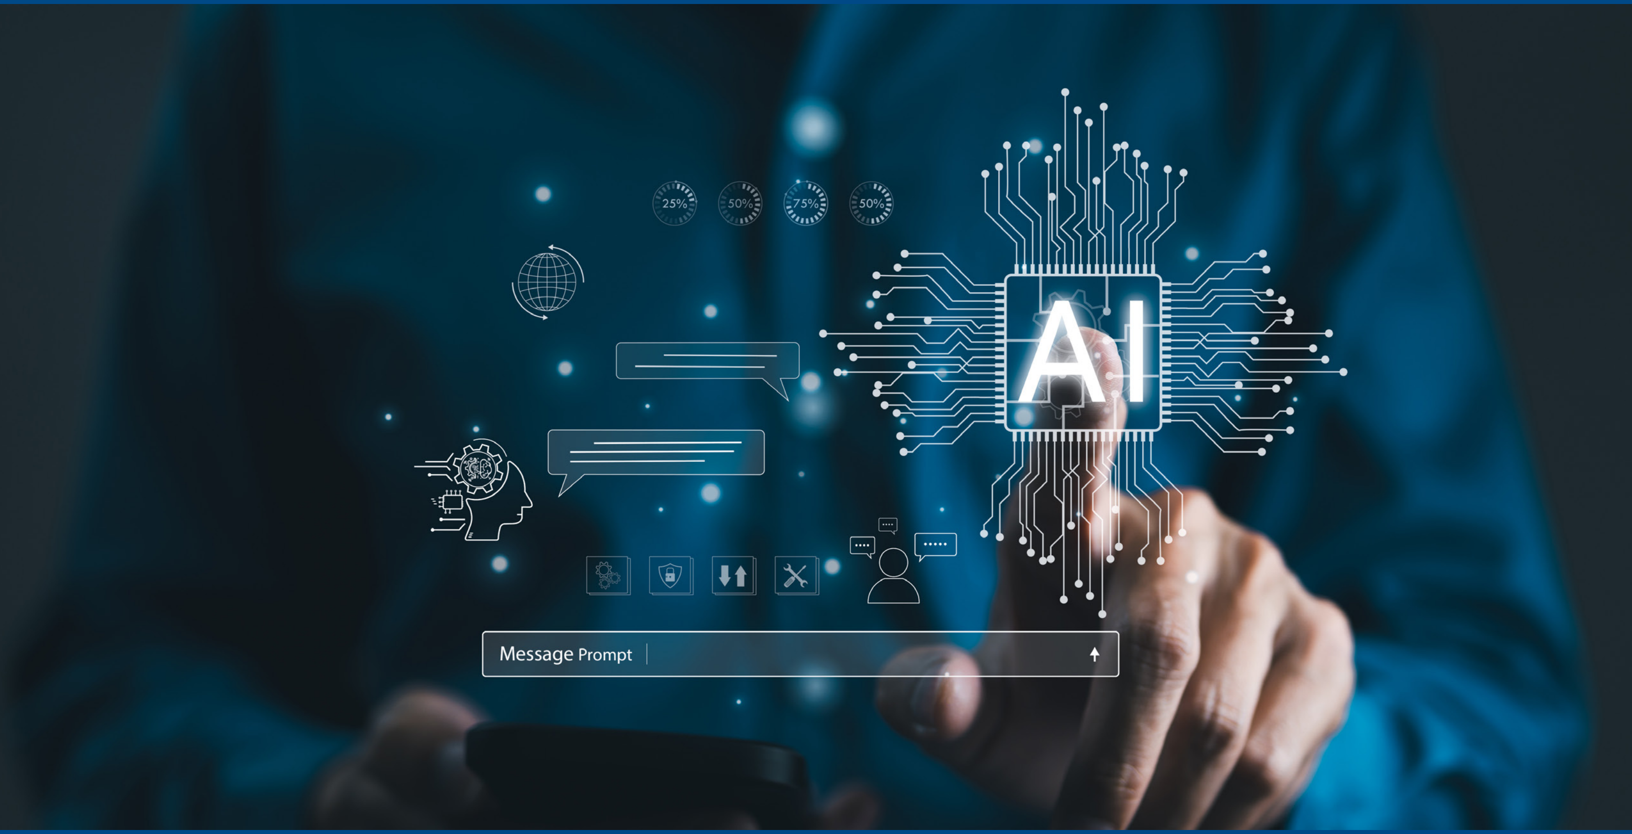

Use of Artificial Intelligence in Financial Aid Offices

Findings From a Review of Institutional AI Policies

NASFAA Research Department

NASFAA Use of AI in the Financial Aid Office Task Force

Table of Contents

Acknowledgements

For this work, NASFAA wishes to thank the members of the Use of AI in the Financial Aid Office Task Force, including:

- Task force chair, Samantha Hicks, Coastal Carolina University
- NASFAA Board of Directors Observer, Celena Tulloss, Auburn University

Task force members:

- Robin Adamson, Coffeyville Community College
- Carly Blagg, Southeast Missouri State University
- Gena Boling, Cornell University
- Crystal Crank, Butler Community College
- Sarah Everitt, Gonzaga University
- Noemi Garcia Tagorda, University of Southern California
- Delia Gomez, Texas A&M University-San Antonio
- Andrew Harvey, Garrett College
- Wayne Kruger, St. Petersburg College
- Jarod Paulson, Capella University
- JoEllen Price, San Jacinto College
- Shara Solero, University of San Diego

Last, we thank the NASFAA members who completed our larger survey on the use of AI in the financial aid office and those who shared their institution's AI policy or guidance. Without your time and feedback, we would not have been able to provide the valuable analysis included in this report to the community.

Definition of Artificial Intelligence

For the purposes of this report, and all work from NASFAA's Use of AI in the Financial Aid Office Task Force, the following definition of AI is used: Artificial Intelligence (AI) is software that simulates human intelligence by interpreting language, recognizing patterns, and generating or analyzing data.

Executive Summary

In late 2025, the National Association of Student Financial Aid Administrators (NASFAA) convened the Task Force on the Use of Artificial Intelligence in Financial Aid Offices to examine how the profession is engaging with artificial intelligence (AI) and to develop recommendations for how NASFAA can best support members as this technology evolves. The task force's work is organized in three phases: a review of relevant literature and reports from peer associations; a national member survey that included an institutional policy scan and listening sessions with financial aid professionals; and a final report to the NASFAA Board of Directors with recommendations. This report presents findings from a review of institutional policies and guidelines uploaded during our membership survey. It is the second report in this task force's series, following Findings From a Survey of Financial Aid Professionals. A third report, Findings From NASFAA Member Listening Sessions, will be published later in 2026. The task force will use the information gathered across these three reports to inform their final report, which will be submitted to the NASFAA Board of Directors. The final report will be published after it has been reviewed and approved by the NASFAA Board in the summer of 2026.

As seen in Findings From a Survey of Financial Aid Professionals, among the more than 800 institutions that participated, only 9% were aware of any policy on AI use, and 27% didn't know if one existed. By comparison, 54% of higher education professionals in other offices reported awareness of policy. To understand what guidance does exist, NASFAA invited survey respondents who indicated their institution had an AI policy to upload it. The documents collected through this process, and the analysis that follows, are best understood in light of what institutional AI policies typically are and do. Institutional policies are, by design, broad in scope — they establish principles, define acceptable use, and set compliance expectations for an entire campus community. They are not written for individual offices or functional areas, and they rarely name specific roles below broad employee categories such as "staff" or "faculty." That is not a limitation unique to AI policy; it reflects how institutional policy works across higher education. The question this scan examines, then, is not whether institutions have tailored their AI policies to financial aid — they have not, and we would not expect them to — but what those broad policies do and do not address, and what that means for financial aid practitioners trying to apply general guidance to a highly specific regulatory context.

About the Document Set

NASFAA received 54 documents from 43 institutions. Upon review, these documents were not uniformly formal, governance-adopted policies. The set includes formal institutional policies (41%), institutional guidelines and guidance documents (35%), strategic frameworks or vision statements (7%), academic integrity or conduct provisions (7%), knowledge base or help articles (7%), and one tool-specific terms-of-use document. This distribution is itself a finding: it reflects a profession at an early stage of policy development, where "policy" means different things at different institutions. Findings throughout this report note where document type affects interpretation. Additionally, because this document set was collected through a survey of financial aid professionals who were being asked about the administrative use of AI, it likely skews toward administrative and operational documents. Academic integrity policies or student-facing AI frameworks that exist at these same institutions may not have been submitted, and the absence of such documents from this scan should not be read as their absence at the institution.

1. The Impact of AI on Work in Higher Education

Key Findings

Most documents cover practitioners, but through broad institutional language rather than role-specific guidance.

Across all 54 documents, 89% cover practitioners in some way. However, only 26% do so explicitly — meaning they name specific roles — and even in those cases, the named roles reflect the broad categories typical of institution-wide policy: "staff," "employees," or "administrative staff." This is consistent with how higher education policies generally function. The practical implication for financial aid practitioners is that they must interpret and apply general institutional guidance to a work context — federal aid compliance, FERPA, Federal Tax Information requirements, high-stakes student decisions — that the policy was not written to address.

Formal policies are stronger across every measured dimension — but make up fewer than half of submitted documents.

Only 22 of the 54 submitted documents (41%) are formal, governance-adopted policies. Among those formal policies, coverage of key provisions is notably stronger: all 22 address disclosure and transparency requirements, 95% address data privacy, 95% name specific approved or prohibited tools, 82% address staff training, and 77% require human oversight of AI outputs. Guidance documents and strategic frameworks are considerably weaker on all of these dimensions. The strength of formal policies relative to informal guidance underscores the practical difference for staff between an institution that has undergone a governance process and one that has issued general guidelines.

Internal staff use of AI is more commonly addressed than student-facing applications — though the collection method likely shapes this finding.

Seventy percent of documents address how staff should use AI in their own administrative work, such as drafting communications, data analysis, and reporting. By contrast, only 52% address AI use in student-facing contexts. It is worth noting, however, that because this document set was collected through a survey of financial aid professionals, respondents likely uploaded the policies most relevant to their own administrative work. Academic integrity frameworks or student-facing AI guidelines that exist at these institutions may simply not have been submitted. The gap between internal and student-facing coverage in this sample should be interpreted with that context in mind.

Data privacy and disclosure are the most consistently addressed provisions; human oversight is less universal.

Across all 54 documents, data privacy (83%) and disclosure requirements (80%) are the most common provisions. Approved or prohibited tools (61%), staff training (61%), and student data protection (61%) appear in roughly 3 in 5 documents. Human oversight — requirements for a human to review AI-generated outputs before use or distribution — is addressed in 57% of documents overall and in 77% of formal policies. Given that the member survey found 41% of financial aid offices have no human review process for AI-assisted work, this gap between what policies articulate and what offices have put into practice warrants attention.

Most documents are restriction-focused, which reflects the nature of institutional policy rather than institutional reluctance.

Fifty-seven percent (57%) of documents are primarily risk- or restriction-focused. This is not surprising: the purpose of institutional policy, across any topic, is to establish compliance expectations and define the boundaries of acceptable use. What is more notable is what is less common — only 15% of documents are balanced in their attention to both risk and use guidance, and very few offer affirmative, operational direction on how to use AI well. For practitioners seeking not just guardrails but practical guidance, the compliance orientation of most institutional documents leaves a meaningful space unfilled.

In this sample, larger institutions were less likely to have submitted formal policies.

Among institutions with enrollment over 10,000 FTE in this sample (n=30), only 13% submitted a formal policy; the majority submitted guidelines, knowledge base articles, or frameworks. By contrast, 74% of medium-sized institutions (2,000-9,999 FTE, n=19) submitted formal policies. This pattern may reflect differences in how large institutions structure AI governance — distributing it across IT, legal, and academic affairs rather than consolidating it in a single policy document — or it may reflect where different institutions are in their policy development timelines. Given the small sample, this finding should be interpreted cautiously.

How most of these policies were developed is unknown.

Fifty-nine percent (59%) of documents provide no information about how they were created or who approved them. Among the 41% that do, most were issued by executive or administrative authority (20%) or through a hybrid process involving a working group and executive approval (17%). Only 4% originated from committee or shared governance processes. The opacity of policy development processes is worth noting: for financial aid staff trying to understand what authority a document carries and whether it reflects input from across the institution, the absence of this information is itself a gap.

Institution-wide AI policies, by design, do not address the specific regulatory context of financial aid work.

Taken together, the documents in this scan reflect institution-wide governance frameworks that were not designed for any specific office or functional area. That is what institutional policy is. What this means in practice for financial aid administrators is that the guidance available to them stops at the institutional level and does not include: broad principles around data privacy, human oversight, and disclosure, applied to a work context defined by FERPA, Federal Tax Information requirements, federal aid compliance, and high-stakes determinations affecting students' access to higher education. The translation of general institutional guidance into operational practice for financial aid offices is work that, in almost every case, has not yet been done.

Policy Scan Findings

What Institutions Are Addressing

Across the 54 documents reviewed, a consistent set of provisions appears most frequently. Data privacy and security (83%) and disclosure or transparency requirements (80%) are the most common, appearing in more than 4 in 5 documents. Approved or prohibited tools, staff training, and student data protection each appear in 61% of documents, and human oversight requirements in 57%. As shown in the table below, formal policies are substantially stronger across all dimensions, particularly in disclosure (100%), data privacy (95%), and approved or prohibited tools (95%).

Provision	All Documents (n=54)	Formal Policies Only (n=22)
Data Privacy / Security	45/54 (83%)	21/22 (95%)
Disclosure / Transparency	43/54 (80%)	22/22 (100%)
Approved or Prohibited Tools	33/54 (61%)	21/22 (95%)
Staff Training	33/54 (61%)	18/22 (82%)
Student Data Protection	33/54 (61%)	18/22 (82%)
Human Oversight of AI Outputs	31/54 (57%)	17/22 (77%)

The gap between formal policies and other document types is consistent and meaningful. Among institutional guidelines, knowledge base articles, and strategic frameworks, provisions appear far less frequently and are typically less specific. Human oversight, for example, is addressed in 77% of formal policies but in only 53% of guidelines and guidance documents. Staff training appears in 82% of formal policies but in 58% of guidelines. This pattern holds across every provision measured, suggesting that the form a document takes, whether it has gone through a governance process or was issued informally, has practical consequences for what it actually requires of staff.

Data Privacy and Security

Data privacy is the most consistently addressed provision across the document set, appearing in 83% of all documents and 95% of formal policies. In practice, however, the guidance provided varies considerably in specificity. The most common form is a prohibition on entering sensitive, confidential, or personally identifiable information into AI tools, a principle that appears in nearly every document addressing data privacy. A smaller share of documents goes further, specifying data classification tiers (e.g., public, confidential, restricted) and linking those tiers to specific tool permissions. A few documents explicitly reference FERPA obligations in the context of AI use, though none specifically address Federal Tax Information (FTI) requirements, which carry distinct handling obligations for financial aid offices.

Disclosure and Transparency

Disclosure requirements appear in 80% of all documents and in all 22 formal policies, making transparency the most universally addressed provision among formal policy documents. The nature of those requirements varies considerably, however. Most disclosure provisions focus on academic contexts, requiring students to disclose AI use in coursework or requiring faculty to disclose when course materials were generated by AI. Fewer documents address disclosure in administrative contexts: whether staff should note when a communication, report, or work product was AI-assisted, and to whom. The administrative disclosure gap is relevant for financial aid offices, where AI-assisted correspondence with students, policy documents, or case notes may raise distinct questions about disclosure that institution-wide academic integrity framing does not resolve.

Approved and Prohibited Tools

Sixty-one percent (61%) of all documents and 95% of formal policies address approved or prohibited AI tools in some way. The most common approach is to designate institutionally supported tools, or an institution-specific deployment, as the preferred or required option, with general caution about using public or consumer-facing AI tools with institutional data. A smaller number of documents maintain explicit prohibited lists, naming specific tools or tool categories that may not be used for institutional purposes. The presence of a tool list in a formal policy is a meaningful signal of institutional governance maturity: it suggests the institution has moved from principle-setting to operational guidance. For staff seeking clarity on what tools are sanctioned for their work, documents that do not address tools at all leave a practical gap.

Human Oversight

Human oversight, defined as the requirement that a person review AI-generated content before it is used or distributed, appears in 57% of all documents and 77% of formal policies. Among the documents that do address it, the guidance ranges significantly. Some documents state a general principle (AI should support, not replace, human judgment) without specifying what review in practice looks like. Others are more operational: one formal policy requires supervisor approval before any AI-assisted content is disseminated publicly; another specifies that AI output should never be used "as is" and requires careful review for accuracy and bias before use. A small number of documents require oversight specifically for consequential decisions, such as personnel actions, aid determinations, and disciplinary matters, distinguishing between lower- and higher-stakes applications. The 41% of documents that do not address oversight at all are predominantly guidelines and informal frameworks, though five formal policies are also silent on the question.

Staff Training

Staff training is addressed in 61% of all documents and 82% of formal policies. As with other provisions, depth varies. The most common references are aspirational, acknowledging that training will be provided or encouraging staff to stay informed, without specifying the content, frequency, or what "trained" means in practice. A smaller number of documents are more specific: one formal policy requires completion of applicable training as a condition of AI use; another mandates annual AI training for all faculty and staff. Several documents note that training will be institution-wide rather than role-specific, which is consistent with the survey finding[2] that only 23% of financial aid professionals who had received any AI training described it as tailored to financial aid or enrollment management.

Student Data Protection

Student data protection provisions that specifically address the handling of student information in AI systems, beyond general data privacy language, appear in 61% of all documents and 82% of formal policies. Most provisions in this category cite FERPA as the governing framework and prohibit the entry of student records or personally identifiable information into AI tools without appropriate data use agreements. A smaller share of documents addresses

2. Findings From a Survey of Financial Aid Professionals

AI-generated outputs that may contain or infer student information, or the use of AI in processes that directly affect student outcomes. Federal Tax Information requirements are not addressed in any document in this scan.

Orientation and Granularity

Fifty-seven percent (57%) of documents are primarily risk- or restriction-focused; 24% are primarily encouragement- or use-focused; and 15% are balanced. As noted in the Executive Summary, the compliance orientation of most documents reflects the nature of institutional policy rather than institutional reluctance toward AI. Notably, no formal policy in this scan is primarily encouragement-focused; all 22 are either risk/restriction-focused (14) or balanced (8). The four strategic frameworks, by contrast, are uniformly encouragement-focused, reflecting their different purpose as vision documents rather than governance instruments.

On granularity, the distribution is relatively even: 35% of documents are mixed (high-level framing with at least some specific actionable provisions), 33% are high-level or general only, and 31% are granular or specific. Among formal policies, none are high-level only; all 22 are either mixed (11) or granular (11), reflecting the governance process that produces them. For a financial aid practitioner seeking guidance they can act on today, a granular or mixed formal policy offers the most usable foundation, while high-level guidance documents offer principles without operational direction.

Tool Approval and Vendor Oversight

Among the 54 documents, 49 specify some form of minimum data security requirements for AI tools, the most consistent governance signal in the entire document set. But the guidance that follows from those requirements is considerably thinner. Only 27 documents name an approval authority for AI tools, with IT offices most commonly designated. Fewer than half of the documents provide any guidance on what the approval process actually entails.

On vendor oversight, 25 documents reference contracting or procurement in some way, and 19 address vendor vetting. In most cases, these are brief mentions rather than operational guidance. Only one document in the sample provides a detailed vendor review process. Three documents describe specific procurement pathways, naming institutional portals or review processes staff can use. The remaining references leave practitioners aware that a process exists, but do not show them how to navigate it.

Tool naming follows a similar pattern. Twelve documents name specific approved tools, and four explicitly prohibit specific tools or tool categories. The remaining 38 mention AI tools in general terms without approving or prohibiting anything specific. Only six documents require an institution to maintain an AI tool inventory or registry, meaning most institutions have no formal mechanism for tracking what tools are actually in use.

The practical implication for financial aid staff is direct: most of these documents would tell a practitioner to be careful with data and to use approved tools, but would not tell them who to contact, what review process to follow, or how long approval might take. The governance infrastructure that would make institutional policy actionable is present in only a small share of the documents reviewed.

Are Practitioners Covered?

The central question motivating this policy scan is whether institutional AI policies address the staff members NASFAA represents: financial aid administrators. Coverage was coded across three categories: explicit (specific non-faculty roles or employee categories named), implicit (practitioners covered under broad "all employees" or "all community members" language, but not named specifically), and not addressed (document scope excludes staff entirely).

Across the 54 documents, 89% cover practitioners in some way — 14 (26%) explicitly and 34 (63%) implicitly. Six documents (11%) do not address practitioners at all; five of those six are either academic integrity provisions or faculty-focused guidelines, documents that were scoped to students or instructional staff from the outset.

Explicit Coverage

Among the 14 documents with explicit practitioner coverage, named roles reflect the broad employee categories typical of institution-wide policy. The most common designations are "staff," "employees," "contractors," and "administrative staff." Two documents specifically name unit heads or administrative leadership. No document in the explicit coverage group names financial aid administrators, enrollment management staff, advisors, or any equivalent student services role. This is consistent with how institution-wide policies function: they establish requirements for broad categories of employees rather than individual offices. However, it means that even the most explicitly staff-oriented documents in this scan do not address the specific regulatory context that financial aid practitioners navigate.

The guidance provided to explicitly covered practitioners follows a consistent pattern across documents. The most common elements are:

- Use only institutionally approved or licensed AI tools
- Do not enter sensitive, confidential, personally identifiable, or restricted data into AI tools
- Review AI-generated outputs for accuracy and bias before use or distribution
- Disclose or acknowledge AI use in work products where required
- Complete required or recommended AI training

A smaller number of documents adds operational specificity: one ties tool permissions to data classification tiers, allowing enterprise-licensed tools to be used with confidential data but not with restricted data; another requires supervisor approval before disseminating AI-assisted content; a third specifies that AI may inform but not serve as the sole basis for personnel, award, or disciplinary decisions.

Implicit Coverage

The 34 documents with implicit practitioner coverage reach staff through a broad institutional scope, using language such as "all community members," "all university employees," or "all individuals granted access to university information." In practice, the guidance these documents provide to practitioners is functionally similar to what explicitly named staff receive: data privacy expectations, disclosure requirements, and accuracy review. The difference is not in content but in legibility. A financial aid staff member reading a policy addressed to "all employees" must infer their own applicability; a policy that names their role removes that interpretive burden.

Among implicitly covered documents, a recurring pattern is that the most substantive guidance is embedded in sections oriented toward faculty or academic use, with administrative staff coverage implied by the opening scope language but not further developed. In several documents, the practitioner-relevant content amounts to a data privacy reminder and a disclosure expectation, with no further guidance on how those principles apply to administrative work.

What Practitioners Are Not Told

Regardless of whether coverage is explicit or implicit, several questions directly relevant to financial aid practitioners are not answered in any document in this scan. None of the 54 documents addresses:

- How FERPA obligations interact specifically with AI tool selection or data entry in administrative processing contexts
- FTI requirements and what they mean for AI use involving tax data
- Whether and how AI may be used in aid eligibility determinations, professional judgment decisions, or Satisfactory Academic Progress reviews
- How to handle AI-generated student communications in terms of disclosure, accuracy review, and institutional accountability
- What to do when an AI tool produces information about a student's aid eligibility that conflicts with staff knowledge or federal guidance

These gaps are not a critique of the documents themselves. Institution-wide policies are not designed to anticipate the operational specifics of individual offices. They do, however, provide a clear picture of where general institutional guidance ends and where financial aid-specific guidance begins.

Patterns by Selected Demographics

Because this document set is relatively small and was collected through a self-selected survey process, patterns by institution type should be interpreted cautiously. Differences across sectors, institution sizes, and regions may reflect genuine variation in how institutions are approaching AI governance, but they may also reflect who responded to the survey, what documents they chose to upload, or where institutions happen to be in their policy development timelines. With that caveat in mind, several patterns are worth noting.

Sector

Sector	n	Formal Policy	Explicit Practitioner Coverage	Data Privacy Addressed	Staff Training Addressed	Human Oversight Addressed
Public 4-Year	25	32%	6 (24%)	92%	68%	64%
Community College	10	50%	2 (20%)	60%	50%	50%
Private Nonprofit	14	43%	3 (21%)	79%	57%	43%
Proprietary	2	50%	1 (50%)	100%	50%	50%
Graduate/Professional	2	50%	1 (50%)	100%	50%	100%

Community colleges stand out in two dimensions in this sample. They are the most likely sector to have submitted a formal policy (50% of community college documents are formal policies) and among the most likely to have balanced or encouragement-oriented documents (30% balanced, compared to 0% among nonprofits). At the same time, community colleges show the lowest rates of data privacy coverage (60%) and explicit practitioner coverage (20%) in the sample — suggesting that while formal policy structures exist, the content of those policies may be less developed than at other institution types. Given the member survey finding that community colleges are among the institutions least likely to provide AI tools to financial aid staff, this pattern may reflect a governance environment that is still in its early stages.

Private nonprofits show the strongest risk/restriction orientation in the sample (79% of nonprofit documents are primarily restrictive), along with relatively high rates of disclosure (86%) and data privacy coverage (79%). Public four-year institutions show the broadest spread across provision types and the largest absolute number of granular or specific documents, which likely reflects their size and the greater resources available for policy development.

Institution Size

The most notable size-related pattern in this sample involves the relationship between enrollment and document type. Among institutions with 10,000 or more Integrated Postsecondary Education Data System (IPEDS) full-time-equivalent (FTE) enrollment, only 13% of submitted documents are formal policies; the majority are guidelines, knowledge base articles, or frameworks. Among institutions with 2,000-9,999 IPEDS FTE, 74% of submitted documents are formal policies. Institutions with fewer than 2,000 IPEDS FTE submitted formal policies in 75% of cases, though the small group size limits interpretation.

This counterintuitive pattern, with larger institutions less likely to have formal policies in this sample, may have several explanations. Large institutions may govern AI through distributed, decentralized frameworks (separate IT policies, academic affairs guidelines, legal and compliance guidance) rather than a single institutional document, and survey respondents at those institutions may have uploaded whichever document was most accessible rather than a primary policy. It may also reflect timing: large institutions with complex governance processes may still be working through policy development.

Region

Regional patterns in this sample are difficult to interpret with confidence, given the small number of documents per region (ranging from five to 14) and the self-selected nature of the document collection.

Methodology

Research Question

This analysis was guided by a single primary research question: How do institutional AI policies address student services practitioners, particularly financial aid administrators? When practitioners are addressed, what guidance or restrictions apply to them?

This question was defined in advance based on NASFAA's specific interest in how existing institutional policy frameworks address the work of financial aid professionals, rather than faculty, students, or the institution as a whole.

Document Collection and De-Identification

As part of NASFAA's national member survey on AI use in financial aid offices, respondents who indicated that their institution had an AI policy were invited to upload it. Documents were submitted voluntarily by survey respondents between January and February 2026. NASFAA received 54 documents representing 43 unique institutions.

Before any analysis was conducted, the NASFAA Research Department reviewed all submitted documents and removed or redacted any information that could identify the submitting institution, including institution names, logos, URLs, and other identifying details. Each document was then assigned a numeric file identifier (e.g., 001.pdf, 002.pdf). A separate key file, maintained exclusively by the Research Department and not used at any stage of analysis, maps those numeric identifiers to institution names.

Separately, NASFAA's membership records were used to compile a demographics file linking each numeric institution code to publicly available institutional characteristics: enrollment (IPEDS FTE), NASFAA member type, sector, and regional association. This file contains no institution names, only numeric codes and demographic variables, and was used solely for the institution-type analysis described in this report.

Analytic Approach

The analysis used an abductive coding strategy combining deductive and inductive elements. The primary deductive dimension, practitioner coverage, was defined in advance in line with the research question. A supporting set of coding variables was also defined before analysis began, covering document type, scope, key policy provisions, orientation, granularity, creation process, and key evidence passages.

Each of the 54 documents was coded on the following variables:

- **Document Type:** Classification as a formal policy, institutional guidelines or guidance, strategic framework, academic integrity or conduct provision, knowledge base or help article, or tool-specific terms of use, with a brief rationale citing specific features of the document.
- **Date and Creation Process:** Effective or adoption date, where stated, and how the document was developed: committee or shared governance, executive or administrative issuance, hybrid process, or not stated.
- **Practitioner Coverage:** Whether the document addresses non-faculty student services staff, coded as explicit (specific roles named), implicit (covered under broad "all employees" or "all community members" language), or not addressed, with a rationale citing specific policy language.
- **Practitioner Guidance:** Where practitioners are covered, what the policy tells them they can or cannot do.

- **Key Provisions:** Whether each of the following is addressed, with a descriptive summary where present: data privacy and security; disclosure and transparency requirements; approved or prohibited tools; staff training; human oversight of AI outputs; student data protection.
- **Internal Use and Student-Facing Tracks:** Whether the document addresses staff use of AI in their own administrative work (Track A) and/or in student-facing interactions such as advising chatbots or automated communications (Track B).
- **Primary Focus and Scope:** The document's main thematic orientation and the population it formally applies to.
- Risk vs. Encouragement Orientation: Whether the document is primarily risk/restriction-focused, encouragement/use-focused, or balanced, with supporting evidence.
- **Policy Granularity:** Whether the document provides high-level principles only, granular and specific operational guidance, or a mix of both.
- **Key Evidence:** Up to three direct quotations most relevant to practitioner coverage, drawn from the source document, formatted with a brief descriptor and page reference.

Technical Implementation and Data Privacy

All automated coding was conducted using Ollama, an open-source framework that runs AI models locally rather than through a cloud-based service. This means that no document content, prompts, or outputs were transmitted to any external server or third-party platform at any point in the analysis. This approach was selected specifically to comply with NASFAA's internal AI Best Practices policy (approved July 2025), which prohibits uploading confidential or proprietary information to generative AI tools. The combination of document de-identification prior to analysis and local model processing ensured that no identifiable institutional content left NASFAA's controlled data environment at any stage.

Human Review

Automated coding constituted a first pass only. The Large Language Model (LLM) output was treated as an initial draft subject to human verification, not as a final analytic product. Consistent with NASFAA's AI Best Practices policy, no unedited AI-generated content appears in this report.

The NASFAA Research Department reviewed the full set of coded documents, reading each source document against its model-assigned codes to assess consistency. The model flagged two documents in the Analyst Notes field as requiring human review due to ambiguous coding decisions; both were examined against the source documents and resolved by the Research Department before analysis proceeded. Where the model's coding was found to be inconsistent with the Research Department's reading of the source document, codes were corrected. All quantitative findings reported in this report are based on those human-reviewed and, where necessary, human-corrected values.

Narrative summaries and interpretive text in the Policy Scan Findings sections were drafted by the Research Department. Generative AI tools supported the initial drafting of portions of this report; all AI-assisted content was reviewed and edited by the Research Department and the task force prior to publication.

Limitations

Several limitations affect the interpretation of the findings in this report.

The document set was self-selected. Documents were uploaded voluntarily by survey respondents who identified their institution as having an AI policy, and the sample captures only what respondents chose to submit and what they understood to constitute a policy. As described throughout this report, the resulting document set is heterogeneous and likely skews toward administrative and operational documents. Academic integrity frameworks, student-facing AI guidelines, or other policy documents at these institutions may not have been submitted.

The sample is small relative to NASFAA's membership. Forty-three institutions are a meaningful starting point, but findings by sector, enrollment size, and region are based on small subgroups and should be interpreted with appropriate caution. Patterns described in the Patterns by Institution Type section are observations rather than generalizable conclusions.

Some documents may be partial representations of a broader governance framework. An institution whose full AI governance spans multiple documents may be represented here by only one of them, which could understate the comprehensiveness of its actual approach.

Finally, the AI policy landscape is evolving rapidly. The documents in this sample were collected in January and February 2026, and institutional policies are currently being revised and expanded. This report is best understood as a snapshot of a moment in time, and the field will continue to develop in the months following publication.