BRODY | GAPP LLP

Clarity in Compliance. Confidence in Litigation.

AI in the Workplace: Acceptable Use Policies, Data Risk, and the Discovery Trap

Prepared by James W. Brody, Esq. | Founding & Managing Partner
James@BrodyGapp.com | 415.246.3995

EXECUTIVE SUMMARY

Every mortgage banking executive, general counsel, and HR leader needs to answer a single question: ***What happens when your employees' AI chat logs become the other side's Exhibit A?***

On February 10, 2026, Judge Jed S. Rakoff of the Southern District of New York ruled in United States v. Heppner that documents a defendant generated using a consumer AI chatbot were not protected by attorney-client privilege or the work product doctrine. The ruling confirmed what compliance professionals should have anticipated: AI conversations are discoverable, they are not confidential on consumer platforms, and they can waive privilege over the very information fed into them. OpenAI's own CEO, Sam Altman, has publicly acknowledged that ChatGPT conversations are not legally protected and can be used as evidence in court.

This handout provides the legal framework, policy architecture, sample provisions, and actionable checklists that mortgage banking institutions need to govern employee AI use—while protecting NPI, preserving privilege, and preparing for a litigation environment where AI chat logs are the new email.

HOW TO USE THIS HANDOUT

General Counsel / Chief Legal Officer: Start with Sections I and IV (the Heppner analysis and the Discovery Trap). Then review Section III (consumer vs. enterprise distinction) and Section VII (discovery preparedness playbook).

Chief Compliance Officer / Risk Officer: Focus on Section II (full regulatory framework), Section V (mortgage-specific risk map), and Section VI (policy architecture and checklists). The 30/60/90-day roadmap in Section VIII provides your implementation timeline.

CISO / IT Leadership: Section III (consumer vs. enterprise platform comparison), Section VI.B (technical controls, DLP, network-level restrictions, BYOD policy), and the ESI inventory checklist in Section VII.

HR Director / CHRO: Section II.B (state employment AI laws) and Section V (the HR row of the risk map) address your specific obligations. The policy checklist in Section VI.B includes employment-specific elements.

I. WHY THIS MATTERS NOW: THE CONVERGENCE OF THREE RISKS

Mortgage banking institutions face a convergence of regulatory, litigation, and operational risk from unmanaged AI use. Three developments have made this an enterprise-level priority, not an IT curiosity.

1

The Heppner Decision Made AI Discovery Real

United States v. Heppner, No. 25-cr-00503-JSR (S.D.N.Y. Feb. 10, 2026) is the first federal ruling to hold that AI-generated documents are not privileged. Judge Rakoff identified three independent grounds: (a) an AI tool is not an attorney and cannot form an attorney-client relationship; (b) consumer AI platform terms of service disclaim confidentiality and reserve rights to disclose data to governmental authorities; and (c) pre-existing documents do not become privileged merely by being transmitted to counsel after the fact. The written opinion, filed February 17, 2026, further suggested that sharing privileged information with a consumer AI tool may constitute a waiver of privilege over the underlying attorney-client communications themselves. Even OpenAI's CEO Sam Altman has publicly acknowledged that ChatGPT conversations are not legally protected and can be used as evidence, describing the issue as one requiring urgent attention.

2

AI Chat Logs Are Electronically Stored Information

Under FRCP Rule 34(a)—and within the scope and proportionality limits of Rule 26(b)(1)—AI prompts, outputs, and activity logs fall squarely within the definition of ESI. In In re OpenAI, Inc., Copyright Infringement Litigation, No. 25-MD-3143 (S.D.N.Y. 2025), Magistrate Judge Ona Wang compelled production of millions of generative AI logs, including user prompts and model responses. The court issued a preservation order directing OpenAI to preserve and segregate output log data that would otherwise be deleted—including user-requested deletions. Courts are not creating special AI exemptions. They are applying settled discovery principles to a new category of records.

3

Mortgage-Specific NPI Exposure Is Acute

Mortgage banking employees routinely handle loan applications, income verification data, credit reports, Social Security numbers, and account information—all of which constitute Nonpublic Personal Information under the Gramm-Leach-Bliley Act. When an employee pastes borrower data into a consumer AI tool, that data leaves the institution's security perimeter and enters a third-party platform that may retain, train on, and potentially disclose it. Beyond the privacy violation itself, inadvertent NPI disclosure to a consumer AI platform may trigger state data breach notification obligations, depending on the jurisdiction and the nature of the data exposed.

II. THE LEGAL FRAMEWORK: FEDERAL & STATE LAWS THAT APPLY

No single "AI law" governs workplace AI use in mortgage banking. Instead, a layered framework of existing federal statutes, new state legislation, and agency guidance creates overlapping compliance obligations. Each layer carries its own enforcement mechanism, and none of them excuse ignorance of AI-specific risks.

A. Federal Framework

Federal Laws and Regulations Applicable to AI Use in Mortgage Banking	
GLBA / Reg. P	Prohibits unauthorized disclosure of NPI to third parties, including AI platforms that may retain, train on, or disclose consumer data. The FTC Safeguards Rule requires written information security programs and risk assessments. 15 U.S.C. §6801–§6809.
ECOA / Reg. B	Requires specific, accurate adverse action reasons reflecting actual decisional factors. AI models used in credit decisions must produce explainable outputs. 15 U.S.C. §1691; 12 C.F.R. Part 1002.
Fair Housing Act	Prohibits discriminatory effects in residential lending regardless of intent. AI tools that produce disparate impact on protected classes create FHA liability. 42 U.S.C. §3601 et seq.
FCRA	Governs access to and use of consumer credit information. AI tools that process or summarize credit report data must comply with permissible purpose and accuracy requirements. 15 U.S.C. §1681 et seq.
FDCPA	Restricts when and to whom debt information can be disclosed. AI tools used in servicing that process debtor information implicate FDCPA privacy provisions. 15 U.S.C. §1692 et seq.
FRCP Rules 26, 34, 37(e)	AI prompts, outputs, and logs constitute ESI under Rule 34(a), discoverable within the scope and proportionality limits of Rule 26(b)(1). Failure to preserve relevant AI data when litigation is reasonably anticipated may result in curative measures under Rule 37(e)(1), or severe sanctions under Rule 37(e)(2) upon a showing of intent to deprive.
SR 11-7	The Federal Reserve's model risk management guidance—the leading U.S. banking supervisory framework for AI/ML governance—is used as a de facto model governance benchmark across financial institutions, including non-bank mortgage lenders, requiring independent validation, ongoing monitoring, and governance documentation.
UDAAP / UDAP	AI-generated communications to borrowers—including marketing, disclosures, and servicing letters—that are misleading, inaccurate, or deceptive create unfair and deceptive acts or practices exposure.

ABA Formal Op. 512	Issued July 29, 2024. Establishes that attorneys using generative AI must comply with duties of competence (Rule 1.1), confidentiality (Rule 1.6), communication (Rule 1.4), and reasonable fees (Rule 1.5). Supervisory lawyers must establish firm-wide AI policies.

B. State AI Laws: The Emerging Patchwork

While the federal government has signaled a preference for a "minimally burdensome" national framework (Executive Order, December 2025), the state-level regulatory landscape is moving faster—and in the opposite direction. A federal moratorium on state AI laws was proposed and then removed from the budget bill (HR 1) before passage. The patchwork is enforceable and growing.

Jurisdiction	Key Obligations	Effective Date / Status
Colorado AI Act (SB 24-205)	Mandatory impact assessments for high-risk AI systems (including employment and lending decisions); risk management programs; public disclosures; deployer governance. AG primary enforcement authority; no express private right of action under the AI Act itself.	June 30, 2026 (delayed from Feb. 1)
Illinois (775 ILCS 5/2-101–110)	Requires disclosure to employees and applicants when AI is used in employment decisions. Covers hiring, promotion, compensation, and termination. IDHR draft implementing rules (Subpart J) would require 4-year recordkeeping for AI notices and records of AI use; confirm final codification status.	January 1, 2026
California FEHA ADS Regulations	Prohibits use of Automated Decision Systems that discriminate based on protected traits. Anti-bias testing (or the absence of it) is explicitly treated as relevant evidence in discrimination claims. 4-year data retention, human oversight, accommodations required. Vendor liability under agency principles.	October 1, 2025

California CPPA ADMT Regulations	Opt-out rights and advance notice when automated decision-making technology replaces human decisions in employment. Applies to CCPA-covered businesses (>$25M revenue). Regulations effective January 1, 2026; ADMT-specific compliance requirements phase in on a later timeline.	January 1, 2026 (ADMT phased later)
Texas TRAIGA (HB 149)	Intent-based discrimination standard; disparate impact alone is insufficient to demonstrate intent. Sandbox program for AI testing. AG exclusive enforcement; no private right of action under this chapter. Minimal private employer obligations; focus on government use.	January 1, 2026
New York City (Local Law 144)	Bias audits for automated employment decision tools used in hiring and promotion within NYC. Annual audit publication. Candidate notification.	In effect (July 2023)
New Jersey Disparate Impact Regs.	Applies LAD disparate impact framework to ADS used in employment decisions; clarifies that automated tools can replicate and amplify existing workforce imbalances.	December 15, 2025

Mortgage banking executives should assume that multi-state operations expose them to the most restrictive applicable standard. A California-originated loan processed by an employee in Colorado using AI tools triggers obligations under both jurisdictions.

III. THE CRITICAL DISTINCTION: CONSUMER AI VS. ENTERPRISE AI

THIS IS THE SINGLE MOST IMPORTANT POLICY DECISION YOUR INSTITUTION WILL MAKE.

Judge Rakoff's Heppner opinion explicitly noted that enterprise AI tools—which contractually guarantee confidentiality and exclude user data from model training—"may present a different analysis." A $20/month consumer subscription does not buy confidentiality. Opting out of data training does not eliminate the platform's right to disclose data to governmental authorities or in

response to legal process—which is the provision Judge Rakoff relied upon. Only enterprise-tier agreements offer the contractual confidentiality protections that could support a privilege argument.

Factor	Consumer Plans	Enterprise Plans
Data used for training?	Yes, by default. Opting out does not eliminate disclosure rights to government or in response to legal process.	No. Enterprise agreements contractually exclude user data from training.
Confidentiality guarantee?	No. Terms reserve right to disclose to governmental authorities and third parties. This was the dispositive factor in Heppner.	Yes. Contractual confidentiality protections, often with BAA and DPA options.
Discoverable in litigation?	Yes. Heppner confirmed: no reasonable expectation of confidentiality under consumer ToS.	Potentially protected. Enterprise confidentiality terms may support privilege arguments, particularly if use is attorney-directed (see Kovel discussion, Section IV.B).
Admin controls?	None. Individual accounts with no centralized oversight.	SSO, domain capture, audit logs, admin dashboards, retention controls.
Compliance certifications?	Generally, none beyond basic privacy policy.	SOC 2 Type II, ISO 27001, HIPAA (varies by vendor), FedRAMP (select plans).
Cost (per user/month)	Free to ~$20	$30–$60+ (varies by vendor, seat commitments)

Frame this correctly for your board: This is not an "AI policy" issue. It is a data governance, confidentiality, and discovery readiness issue. Treat AI chat logs the way you treat email, Slack, and text messages—retain them, govern them, and be prepared to defend them.

Bottom line for mortgage banking institutions: If your employees are using consumer versions of ChatGPT, Claude, Gemini, or Copilot to process any information related to their work—including borrower data, internal strategy, loan files, or HR matters—your institution has an NPI exposure problem, a potential privilege waiver problem, and an ESI preservation gap. The remediation starts with either (a) prohibiting consumer AI use for work purposes entirely, or (b) deploying an enterprise-grade platform with contractual controls. This analysis applies equally to employees using personal devices—the BYOD gap is where the most uncontrolled AI usage occurs, and institutional policy must extend to personal device use for work-related AI activity.

IV. THE DISCOVERY TRAP: AI CHAT LOGS AS THE NEW EMAIL

Most mortgage banking executives and general counsels have not yet internalized a reality that litigators are already weaponizing, which is that AI chat logs are the richest, most revealing, and most dangerous category of ESI that has emerged since email. And unlike email—which people generally understand is a business record—employees treat AI chatbots like a private conversation. They are anything but.

A. What Makes AI Chat Logs So Dangerous

1

Employees Use AI as a Confessional

Employees paste termination memos into AI tools asking "is this legal?" They upload discrimination complaints asking "how should I respond?" They input performance reviews asking "can this employee sue me?" They are treating a public commercial platform as if it were their personal attorney—and the information they disclose is neither privileged nor confidential.

2

The Content Reveals State of Mind

In litigation, an employee's AI prompts can establish knowledge, intent, motive, and timing more directly than almost any other form of ESI. A supervisor who asks an AI tool "how do I fire someone who just filed a workers' comp claim without getting sued" has created a document that plaintiff's counsel will frame as premeditation. The prompt is dated, detailed, and in the employee's own words.

3

Deletion Is Not Destruction

Courts have ordered AI platforms to preserve and segregate AI logs that would otherwise be deleted, including user-requested deletions. In In re OpenAI (May 13, 2025), Magistrate Judge Wang issued a preservation order requiring OpenAI to retain output log data across consumer ChatGPT accounts—data that OpenAI had been routinely deleting. A user's belief that deleting a conversation eliminates the record is factually wrong—and acting on that belief after a litigation hold triggers could constitute spoliation.

4

The Volume Is Staggering

Unlike email—where a discovery request might target specific custodians over a defined period—AI usage generates massive volumes of interactions across every employee who has access. A single employee's ChatGPT history can contain thousands of prompts spanning months. The proportionality arguments that limit email discovery are still being developed for AI data, and courts are not yet consistently applying them in the defendant's favor.

5

Personal Devices Multiply the Exposure

Employees using personal phones, tablets, or laptops to access consumer AI tools for work-related purposes create ESI outside the institution's control. The institution may have no visibility into these accounts, no ability to implement a litigation hold over them, and no

contractual right to access the data—yet the data may be discoverable if it relates to the employee's work. A BYOD policy that does not address AI tool usage is incomplete.

B. The Privilege Waiver Problem — and the Kovel Pathway

The Heppner decision's most consequential holding may be its footnote: Judge Rakoff observed that sharing privileged information with a consumer AI platform could waive the privilege over the underlying attorney-client communications themselves. This means that if an employee receives legal advice from in-house or outside counsel and then inputs that advice into a consumer AI tool, the privilege over the original counsel communication may be destroyed—not just over the AI-generated output, but over the source communication.

For mortgage banking institutions, the practical scenarios are predictable: An HR director receives advice from employment counsel about a pending termination and then inputs the counsel's analysis into Claude to "summarize the key points." A compliance officer receives a legal memorandum on a RESPA investigation and uploads it to ChatGPT to "identify the main risks." A loan officer receives notice of a repurchase demand and uses an AI tool to "draft a response strategy." In each case, the employee may have destroyed the privilege over the original legal communication—and created a new, discoverable document in the process.

THE KOVEL PATHWAY: HOW TO PRESERVE PRIVILEGE WHILE USING AI

Judge Rakoff's written opinion left the door open. He stated that had counsel directed Heppner to use Claude, the tool "might arguably be said to have functioned in a manner akin to a highly trained professional who may act as a lawyer's agent within the protection of the attorney-client privilege." This references the Kovel doctrine (United States v. Kovel, 296 F.2d 918 (2d Cir. 1961)), which extends privilege to non-lawyer professionals—such as accountants or consultants—retained by counsel to assist in rendering legal advice.

Practical implication for institutions: If an enterprise AI tool is deployed under counsel's direction, under an enterprise agreement that guarantees confidentiality, and used as part of the attorney's workflow for rendering legal advice, there is a plausible argument that the Kovel doctrine extends privilege protection to the AI-assisted work product. The elements to support this argument are: (1) the AI tool is used at counsel's direction, (2) its use is necessary to facilitate counsel's rendering of legal advice, (3) the platform's contractual terms support a reasonable expectation of confidentiality, and (4) the resulting work product reflects counsel's legal strategy and judgment. This is not yet settled doctrine for AI tools; no court has affirmatively extended Kovel to an AI platform, and institutions should treat this as a risk-managed argument supported by the architecture of their AI deployment—not a guarantee. But it provides the structural blueprint for building a defensible AI usage protocol within the attorney-client relationship.

C. Litigation Hold Obligations

Once litigation is reasonably anticipated, the duty to preserve evidence attaches. Under FRCP 37(e), if ESI that should have been preserved is lost because a party failed to take reasonable steps to preserve it and it cannot be restored or replaced, courts may order curative measures no greater than necessary to remedy the prejudice. The most severe sanctions—adverse inference instructions, dismissal, or default judgment—are available only upon a finding that the party acted with the intent to deprive another party of the information's use in the litigation. Negligence and even gross negligence do not support those severe measures under the 2015 amendments. This means that failing to preserve AI chat logs will not automatically result in catastrophic sanctions—but it will result in curative measures, it will damage institutional credibility, and if intent is shown, the consequences are case-ending. Institutions must update their litigation hold procedures, document preservation policies, and custodian identification protocols to include AI platforms/accounts—including accounts on personal devices used for work.

V. MORTGAGE-SPECIFIC AI USE CASE RISK MAP

The following table identifies the primary contexts in which mortgage banking employees interact with AI tools, the specific legal risk each interaction creates, and the policy controls that should be in place.

Use Case / Dept.	What Employees Actually Do	Legal Risk Without Policy Controls
Loan Production / LOs	Paste borrower income docs and credit data into AI to summarize or analyze; draft pre-approval letters; generate marketing content targeting specific borrower profiles.	GLBA/NPI violation; ECOA fair lending risk if AI output influences credit terms; UDAAP if AI-generated marketing is misleading; all content discoverable.
Underwriting / QC	Use AI to review appraisal reports, flag exceptions, or analyze DU/LP findings; draft condition letters.	Model risk under SR 11-7 if AI influences credit decisions; adverse action explainability failures under Reg. B; Freddie Mac Guide §1302.8 governance obligation.
Servicing	AI-assisted loss mitigation analysis; AI-generated borrower communications; default servicing workflow automation	FDCPA disclosure limits; UDAAP deceptive communications; Fannie/Freddie servicing guide compl.; all chat logs and outputs are ESI
Human Resources	Draft termination memos, PIP letters, and severance calculations using AI; screen resumes; analyze compensation data.	Illinois AI disclosure law (eff. 1/1/26); CA FEHA ADS regs (eff. 10/1/25); NYC LL144 bias audit; NJ disparate impact regs (eff. 12/15/25); privilege waiver if legal advice is input into consumer AI.
Legal / Compliance	Summarize regulatory guidance; draft policy documents; analyze enforcement actions; prepare for examinations.	ABA FO 512 competence and confidentiality obligations; privilege waiver risk per Heppner; Kovel doctrine opportunity if structured properly; vendor data training concerns.

Executive / C-Suite	Strategic planning; M&A analysis; competitive intelligence; board reporting.	Trade secret exposure; discoverable evidence of decision-making rationale in shareholder or regulatory disputes; privilege waiver overboard communications if AI is used to process them.
Information Technology	Code generation; system documentation; vendor evaluation; cybersecurity analysis.	Exposure of system architecture, access credentials, or security vulnerabilities to third-party platforms; potential GLBA Safeguards Rule violations.
Whistleblower / Compliance Reporting	Employees use AI to research whether conduct constitutes fraud; draft internal complaints; prepare summaries for regulators or counsel	AI logs may reveal protected whistleblower activity (Dodd-Frank §922, SOX §806, FCA); discoverable evidence of what employee knew and when; may also reveal lack of institutional response

VI. BUILDING THE AI ACCEPTABLE USE POLICY: DECISION FRAMEWORK

An AI Acceptable Use Policy for a mortgage banking institution is not a standalone document. It is an integrated layer within the institution's existing compliance, information security, vendor management, and employment frameworks.

A. Threshold Decision: Enterprise Platform, Prohibition, or Hybrid

Three Policy Models and Their Tradeoffs	
Model 1: Full Prohibition	Prohibit all employee use of generative AI for any work-related purpose. Simplest to enforce but creates shadow IT risk, competitive disadvantage, and employee friction. May be appropriate for small shops without IT infrastructure to manage enterprise deployment.
Model 2: Enterprise-Only	Deploy an enterprise AI platform (e.g., ChatGPT Enterprise, Claude Business/Enterprise, Microsoft Copilot) with admin controls, SSO, audit logging, and data retention policies. Prohibit consumer AI use for work purposes. Most defensible approach for institutions with compliance infrastructure.
Model 3: Hybrid / Tiered	Allow limited consumer AI use for non-sensitive tasks (e.g., formatting, grammar, general research) while requiring enterprise-grade tools for any activity involving NPI, borrower data, legal analysis, HR matters, or confidential business information. Requires clear data classification and employee training.

B. Essential Policy Elements

✓	AI Acceptable Use Policy — Required Components
☐	**Scope:** Define which AI tools are approved, prohibited, and conditionally approved. Include chatbots, coding assistants, transcription tools, image generators, and AI features embedded in existing software.
☐	**Data Classification:** Establish clear categories (Public, Internal, Confidential, NPI/Restricted) and specify which categories may and may not be processed through any AI tool.
☐	**NPI Prohibition:** Explicit prohibition on inputting borrower NPI, credit data, loan file contents, SSNs, income documentation, or any data subject to GLBA, FCRA, or Reg. P into any AI tool not covered by an enterprise agreement with contractual NPI protections.
☐	**Privilege Preservation:** Explicit prohibition on inputting attorney-client privileged communications, legal memoranda, litigation strategy, or information received from counsel into any AI tool—consumer or enterprise—unless the use is directed by counsel under a Kovel-structured protocol.
☐	**Approved Use Cases:** Define specifically what employees may use AI for and what requires approval or is prohibited.
☐	**HR and Employment Decisions:** If AI is used in any employment decision, disclose that use to affected employees per applicable state law (IL, CA, NJ, NYC). Conduct and publish bias audits as required.
☐	**BYOD and Personal Devices:** Extend the AI acceptable use policy to personal devices used for any work-related purpose. Require employees to disclose personal AI accounts used for work. Include personal device AI accounts in litigation hold procedures.
☐	**Vendor Due Diligence:** Require review of any AI vendor's terms of service, privacy policy, data training practices, retention policies, and government disclosure provisions before institutional deployment.
☐	**Monitoring and Audit:** Establish that the institution may monitor AI tool usage on company devices and networks. Implement monitoring with required notices/consents under applicable state electronic monitoring statutes (e.g., NY, CT, DE) and coordinate with employment counsel.
☐	**Incident Response:** Include AI-related data incidents in the institution's existing incident response and breach notification framework. Map AI data incidents to applicable state breach notification statutes.
☐	**Retention and Preservation**: Integrate AI chat data into the institution's document retention schedule. Define retention periods consistent with regulatory requirements.
☐	**Litigation Hold Integration:** Update litigation hold templates and custodian questionnaires to specifically identify AI platform accounts, chat histories, and output files as ESI sources—including personal device accounts.
☐	**Training and Acknowledgment:** Require annual training on AI acceptable use. Require signed acknowledgment from all employees, including executives.

☐	**Disciplinary Consequences:** Specify that violation of the AI policy is subject to the same disciplinary framework as other information security and compliance policy violations.
☐	**GSE Governance Alignment:** For Freddie Mac Seller/Servicers, confirm that the AI governance framework satisfies Guide §§1302.2 and 1302.8 (enterprise-wide effective March 3, 2026). For Fannie Mae, align with Information Security and Business Resiliency Supplement.

C. Sample Policy Provisions

The following illustrative provisions may be adapted for use in an institution's AI Acceptable Use Policy. These are starting points for drafting, not final policy language, and should be reviewed by qualified counsel for jurisdiction-specific requirements.

SAMPLE PROVISION 1 — NPI Prohibition

"No employee, contractor, or agent of [Institution] shall input, upload, paste, or otherwise transmit any Nonpublic Personal Information (as defined by the Gramm-Leach-Bliley Act and Regulation P), consumer credit information (as defined by the Fair Credit Reporting Act), or any data from which a consumer's identity, financial condition, or account information could be derived, into any generative artificial intelligence tool, chatbot, or automated text generation platform that is not an Approved Enterprise AI Platform as designated by the Information Security Officer."

SAMPLE PROVISION 2 — Privilege Preservation

"No employee shall input into any AI tool—including Approved Enterprise AI Platforms—any communication received from in-house counsel, outside counsel, or any attorney acting on behalf of the institution, or any document marked as privileged, attorney work product, or attorney-client communication, unless such use has been specifically authorized in writing by the General Counsel or designated Legal Department representative and is conducted under a protocol designed to preserve applicable privilege protections."

SAMPLE PROVISION 3 — Personal Device / BYOD

"This Policy applies to all AI tool usage for work-related purposes, regardless of whether the tool is accessed on a company-owned device or a personal device. Employees who use personal devices to access AI tools for any work-related purpose must: (a) comply with all provisions of this Policy as if using a company device; (b) disclose all personal AI tool accounts used for work-related purposes to the Information Security Officer; (c) preserve all AI-related data on personal accounts when notified of a litigation hold; and (d) make such accounts available for review upon request by the Legal Department in connection with any investigation, audit, or litigation."

VII. DISCOVERY PREPAREDNESS: THE AI ESI PLAYBOOK

Litigation counsel across the country are adding AI-specific discovery requests to their standard repertoire. Within the next twelve to eighteen months, mortgage banking institutions should expect to receive discovery demands targeting AI usage in employment disputes, fair lending investigations, repurchase litigation, and CFPB enforcement actions.

A. Update Your ESI Inventory

✓	AI-Specific ESI Inventory Checklist
☐	Identify all AI platforms in use (approved and unapproved) across the institution, including embedded AI features in existing tools and personal device usage.
☐	Map data flows: What data goes into the AI tool? What comes out? Where is it stored? Who can access it?
☐	Determine which AI platforms retain user data, for how long, and under what conditions data can be recovered or produced.
☐	Identify custodians with significant AI usage patterns, particularly in litigation-sensitive roles (HR, legal, compliance, executive).
☐	Confirm whether enterprise AI agreements include a right to obtain or export user activity logs for litigation purposes.
☐	Document the institution's AI tools in the same manner as email systems, document management platforms, and communication tools.
☐	Confirm whether AI memory or conversation history features are enabled and whether they can be centrally managed or exported.

B. Update Litigation Hold Procedures

Litigation hold notices must now specifically instruct custodians to preserve AI-related data. A litigation hold that does not mention AI chat logs, prompts, and outputs is incomplete. The hold notice should instruct custodians to (a) not delete any AI chat history, (b) not disable or modify AI account settings that affect data retention, (c) identify all AI platforms on which they have accounts (personal and institutional), (d) preserve any downloaded or exported AI outputs, and (e) not disable AI memory or conversation history features pending further instruction from counsel.

C. Prepare for Offensive and Defensive AI Discovery

Discovery Scenario	What Opposing Counsel Requests	How to Prepare
Employment Dispute (Wrongful termination, discrimination)	AI chat logs of supervisors involved in the decision; prompts related to termination strategy, legal exposure assessment, performance evaluations	Policy prohibiting HR use of consumer AI for employment decisions; privilege protocols; enterprise AI with audit logs and admin retention controls
Fair Lending Investigation	Documentation of AI models used in underwriting, pricing, or marketing; model validation records; AI governance policies; chat logs of employees discussing model behavior	Freddie Mac §1302.8 compliance documentation; model inventory; AI governance framework; bias testing records
Repurchase / GSE Dispute	Evidence of AI involvement in QC, underwriting, or document verification; AI tool outputs that influenced loan decisions	Clear AI use case registry; human review documentation showing AI was assistive, not determinative; GSE-compliant governance records
LO Transition / Trade Secret Litigation	AI chat logs of departing LOs showing pre-departure activity; evidence of proprietary data input into AI tools; competitor intelligence queries	Monitor AI usage for departing employees; include AI platforms in exit procedures; contractual restrictions on AI-assisted competitive activity
Whistleblower / Qui Tam / FCA	AI chat logs showing employee researched fraud or compliance violations; evidence of when employee knew of wrongdoing; AI-generated complaint drafts	Preserve all AI data that may evidence protected activity; ensure compliance reporting channels are documented separately from AI usage
CFPB / State AG Enforcement	AI governance policies; vendor due diligence records; consumer complaint analysis showing AI involvement; training records	Exam-ready AI governance dossier; documented risk assessments; vendor contracts with AI-specific controls; board-level AI reporting

VIII. IMMEDIATE ACTION ITEMS: THE 30/60/90-DAY ROADMAP

DAY 1 — EMERGENCY DIRECTIVE: Issue a same-day written directive to all employees: (1) Do not input any NPI, borrower data, privileged communications, or confidential business information into any consumer AI tool effective immediately. (2) Do not delete any existing AI chat histories. (3) If you have used a consumer AI tool for any work-related purpose, notify [designated officer] within 48 hours. This directive should be issued by the CEO or General Counsel and distributed by email with read-receipt confirmation.

Days 2–30: Foundation

✓	30-Day Priority Actions
☐	Conduct an institutional AI census: identify every AI tool in use (approved and shadow IT), every employee with access, and every data category being processed. Include personal device usage.
☐	Brief the board or executive committee on the Heppner decision and its implications for institutional litigation exposure.
☐	Engage outside counsel to conduct a privilege waiver risk assessment across current AI usage patterns.
☐	Update litigation hold templates to include AI platforms and chat histories as ESI categories.
☐	Evaluate enterprise AI platform options and initiate procurement process.

Days 31–60: Architecture

✓	60-Day Build-Out
☐	Select and deploy an enterprise AI platform with contractual confidentiality, admin controls, audit logging, and SOC 2/ISO compliance.
☐	Draft a comprehensive AI Acceptable Use Policy incorporating all elements identified in Section VI of this handout, including the sample provisions in Section VI.C.
☐	Establish a cross-functional AI governance committee (legal, compliance, IT, HR, operations) with a defined charter and meeting cadence.
☐	Conduct vendor due diligence on all existing AI tools and integrations, evaluating ToS, privacy policies, data training practices, and government disclosure provisions.
☐	Integrate AI governance into the Freddie Mac Guide §§1302.2 and 1302.8 compliance framework.

Days 61–90: Operationalization

✓	90-Day Operational Maturity
☐	Deliver institution-wide AI Acceptable Use training with signed acknowledgment. Include specific scenarios relevant to each department (LO, UW, HR, servicing, legal).
☐	Implement technical controls: block unauthorized AI domains on company networks; deploy DLP tools to detect NPI in outbound AI traffic; enable enterprise AI audit logs; address BYOD monitoring capabilities.
☐	Conduct a tabletop exercise simulating an AI-related discovery demand in an employment or fair lending dispute.
☐	Publish the AI governance framework for GSE examination readiness, including model inventory, risk assessment, and board reporting.
☐	Schedule quarterly AI policy reviews to address new legal developments, platform changes, and emerging case law.

KEY AUTHORITIES AND SOURCES

Case Law

1. United States v. Heppner, No. 25-cr-00503-JSR (S.D.N.Y. Feb. 10, 2026) (bench ruling); Written Op. filed Feb. 17, 2026, 2026 WL 436479.

2. In re OpenAI, Inc., Copyright Infringement Litigation, No. 25-MD-3143 (S.D.N.Y. 2025) (Magistrate Judge Ona Wang, preservation order May 13, 2025; production order 2025 WL 3468036).

3. United States v. Kovel, 296 F.2d 918 (2d Cir. 1961) (extension of attorney-client privilege to non-lawyer agents of counsel).

Federal Statutes and Regulations

4. Gramm-Leach-Bliley Act, 15 U.S.C. §§6801–6809; FTC Safeguards Rule, 16 C.F.R. Part 314; Regulation P, 12 C.F.R. Part 1016.

5. Equal Credit Opportunity Act / Regulation B, 15 U.S.C. §1691; 12 C.F.R. Part 1002.

6. Fair Housing Act, 42 U.S.C. §§3601 et seq.

7. Fair Credit Reporting Act, 15 U.S.C. §1681 et seq.

8. Fair Debt Collection Practices Act, 15 U.S.C. §1692 et seq.

9. Fed. R. Civ. P. 26(b)(1), 34(a), 37(e) (discovery, ESI definition, and preservation of ESI).

10. Federal Reserve SR 11-7: Guidance on Model Risk Management (Apr. 4, 2011).

11. Freddie Mac Bulletin 2025-16, amending Guide Sections 1302.2 and 1302.8 (enterprise-wide AI governance, eff. March 3, 2026).

State Laws

12. Colorado AI Act (SB 24-205), as amended by SB 25B-004 (eff. June 30, 2026).

13. Illinois Human Rights Act, as amended by HB 3773, 775 ILCS 5/2-101–110 (eff. Jan. 1, 2026).

14. California FEHA ADS Regulations, Cal. Code Reg. tit. 2, §§11008–11097 (eff. Oct. 1, 2025).

15. California CPPA ADMT Regulations (eff. Jan. 1, 2026; ADMT compliance phased).

16. Texas Responsible AI Governance Act (TRAIGA), HB 149 (eff. Jan. 1, 2026).

17. New York City Local Law 144 (automated employment decision tools; in effect July 2023).

18. New Jersey Disparate Impact Regulations governing ADS in employment (eff. Dec. 15, 2025).

Professional Guidance and Federal Policy

19. ABA Standing Committee on Ethics and Professional Responsibility, Formal Opinion 512, "Generative Artificial Intelligence Tools" (July 29, 2024).

20. White House Executive Order on AI (Dec. 2025); AI Action Plan (July 2025).

BRODY | GAPP LLP

National Mortgage Banking Regulatory Compliance and Litigation Practice
IMBs · Depositories · Credit Unions · Mortgage Brokers · Fintechs