AIGovXRay

Public AI Governance Signals Across American Banking

2026 — First Edition

Report scope: 35 US federally insured depository institutions with observable AI interaction surfaces across public-facing digital channels, scanned via AIGovXRay Web Surface Scan during Q2 2026.

Generated: June 2026

Scan Basis: Web Surface Scan — observable page signals only. No backend, runtime, or log evidence was analysed.

Issued by: Mechapiens PC | Product: AIGovXRay Web Surface Scan

Evidence Fragility: All findings classified HIGH. Single point-in-time capture; cannot be corroborated by runtime or backend evidence.

1. Scan Methodology

Evidence Scope

All findings in this report derive exclusively from Web Surface Scans. A Web Surface Scan captures the observable page layer of a target URL: rendered HTML, HTTP response headers, DOM snapshot, and interaction log. No backend systems, runtime environments, internal logs, or vendor-internal controls were inspected.

This scope is intentional. A web surface scan produces evidence that any regulator, auditor, or counterparty can independently verify from a public vantage point, without privileged access. It is not a substitute for a runtime audit, a log-based deep scan, or an internal control assessment.

Corpus

The scan corpus comprises 35 entities: US federally insured depository institutions — commercial banks and federally insured credit unions — selected on a single criterion: the presence of an observable AI or chatbot interaction surface on a public-facing digital channel accessible without authentication.

Corpus construction proceeded in two stages: (1) identification of US banking institutions with vendor-confirmed or self-disclosed AI chatbot deployments on their public websites, and (2) AIGovXRay Web Surface Scan of each candidate URL. Institutions whose AI systems operate exclusively behind authentication, within mobile applications, or in back-office environments were excluded. The corpus therefore represents the observable public AI surface layer of the market, not the full US banking market.

Regulatory Framework

Unlike the European edition of this report, this edition operates in an environment where no single enacted federal AI-specific statute with direct deployer obligations was in force at the time of scanning. The Great American AI Act (GAAIA), a bipartisan discussion draft released in June 2026, proposes transparency and accountability requirements for AI deployers but had not been enacted.

Regulatory signals are mapped to four reference frameworks:

- GAAIA — draft disclosure and transparency provisions
- NIST AI RMF — the voluntary risk management framework referenced by federal agencies and major financial regulators
- GLBA / CFPB guidance — applicable to consumer data processed through AI interfaces
- ISO/IEC 42001 — the AI management system standard providing governance and data-entry risk benchmarks

No finding in this report constitutes a determination of non-compliance with any enacted law.

Finding Categories

Category	Severity	Description
Client-side Application Identifier Review	Low	A client-visible application identifier was detected in frontend-delivered code; does not by itself distinguish agent routing from benign SDK, monitoring, or telemetry usage.
Agent Backend Exposure	Medium	A client-visible API endpoint or identifier specifically associated with agent or chatbot proxy infrastructure was detected with sufficient contextual specificity.
AI Disclosure Gap	Medium	A visible AI interaction surface was detected without a co-located user-facing disclosure of AI nature or limitations on the same captured page.
Data Governance Gap	Medium	A visible user data-entry surface connected to an AI interaction context was detected without a co-located privacy, terms, or data-processing policy link.

Evidence Fragility

All findings carry a fragility classification of high. Findings are derived from a single point-in-time page capture and cannot be corroborated by runtime traces, server-side logs, or backend access-control records.

A high-fragility finding is a governance signal requiring follow-up investigation — not a concluded determination of non-compliance.

2. Corpus Overview

Cohort	Entities	Share
Entities with at least one finding	3	8.6%
Entities with zero findings	32	91.4%
Total corpus	35	100%

The 32-entity majority with zero findings should not be read as a uniform governance clean-bill. Absence of findings reflects either that no AI-related surface artefacts were observable at the scanned URL at the time of capture, or that the observable surface met the engine's heuristic thresholds for the finding categories in scope. A surface scan cannot evidence the presence of strong backend controls, nor can it evidence their absence.

Market Structure Note

The American banking market exhibits a pattern not observed in the European edition: a deliberate 'AI-behind-authentication' deployment strategy among the largest institutions. All eight US Global Systemically Important Banks (G-SIBs) — JPMorgan Chase, Bank of America, Wells Fargo, Citigroup, Goldman Sachs, Morgan Stanley, State Street, and BNY Mellon — were assessed and excluded from the scan corpus because their AI systems operate exclusively within authenticated digital banking environments or mobile applications inaccessible to a public-vantage web surface scan.

This is not a finding; it is a corpus boundary determination. The observable AI surface in the US banking market is concentrated disproportionately among community banks and federally insured credit unions, which deploy third-party AI chatbot platforms directly on their public-facing websites. This structural pattern is the inverse of the European market.

3. Findings

3.1 Institutions with Confirmed Signals

Institution	URL	Risk	Finding	Observable Signal
Harvard Federal Credit Union	harvardfcu.org	MEDIUM	Data Governance Gap	File upload affordance without co-located privacy / policy link
Citadel Credit Union	citadelbanking.com	LOW	Client-side App ID Review	ApplicationId header in frontend JS; privacy policy link present
Capital One	capitalone.com	LOW	Client-side App ID Review	applicationID field in frontend JS (New Relic NREUM pattern)

3.2 Finding Frequency Across the Corpus

Finding Category	Institutions Affected	% of Corpus	Notes
Client-side Application Identifier Review	2	5.7%	Unclassified applicationID or SDK header; may be benign telemetry
Agent Backend Exposure	0	0.0%	No confirmed named AI proxy endpoint detected in corpus
AI Disclosure Gap	0	0.0%	No visible AI UI without disclosure detected at scan time
Data Governance Gap	1	2.9%	Visible data-entry surface without co-located policy link (Harvard FCU)

3.3 Client-side Application Identifier Review — Citadel CU & Capital One

Two institutions produced a Client-side Application Identifier Review finding — the lowest-severity category. In both cases the evidence is a generic applicationID field in frontend-delivered HTML or JavaScript, and the individual reports explicitly state that the captured evidence does not by itself distinguish agent routing from benign SDK, monitoring, or telemetry usage.

Capital One: The observed artefact is a New Relic NREUM browser monitoring configuration — a standard web analytics SDK with no AI-specific function. The individual report assigns a low confidence score (0.62) and recommends identifying the owning SDK before treating the signal as backend exposure. Capital One does deploy a named AI assistant, Eno, with a dedicated public page; however, the scan of the capitalone.com homepage did not detect Eno-specific identifiers at scan time.

Citadel Credit Union: The ApplicationId detected in frontend-delivered JavaScript is embedded within an Azure Application Insights instrumentation string — a standard Microsoft telemetry SDK. Citadel does operate a named AI chatbot, Adel, deployed via Posh AI; the Application Insights identifier is a separate, co-present monitoring artefact.

> These findings should prompt internal SDK documentation review, not governance remediation.

3.4 Data Governance Gap — Harvard Federal Credit Union

Harvard Federal Credit Union (harvardfcu.org) is the only institution to produce a finding above the lowest severity tier. The scan assigned an overall risk signal of medium, with a single Data Governance Gap finding at medium severity and a confidence score of 0.60.

What the Scan Observed

The AIGovXRay engine detected a file upload affordance on the scanned page — a visible data-entry surface that accepts user-submitted files — without a co-located privacy, terms, or data-processing policy link on the same captured page. The evidence excerpt recorded by the engine is: 'Upload a Document'. The interaction log, DOM snapshot, HTTP headers, and screenshot were all captured and are referenced in the individual report.

The individual report explicitly states what this evidence cannot prove: what user-facing policy terms govern data submitted through this visible interface, and whether retention, processing, or sharing expectations are disclosed to users in the captured state.

Why This Finding Has AI Governance Relevance

Harvard FCU is a Posh AI customer. Posh AI deploys a public-facing digital assistant — accessible from the institution's homepage without authentication — that handles member queries, FAQ responses, and financial product guidance. The file upload capability detected occurs within or adjacent to this AI-assisted member service surface.

When a user submits a document through an AI-assisted interface without a visible data-processing disclosure, the conditions for a Data Governance Gap finding are met: a visible data-entry surface connected to an AI interaction context, without a co-located policy link.

Comparison with the European Edition

In the European edition, the Data Governance Gap finding was associated with Piraeus Bank and Česká spořitelna — both presenting visible AI chat interfaces with prompt input fields and no policy accompaniment. The Harvard FCU finding presents a structurally similar configuration, with a key difference: the European findings were triggered by a chatbot textarea input, while the Harvard FCU finding is triggered by a file upload control — a higher-sensitivity data-entry modality, as uploaded documents typically contain structured personal or financial information.

The NIST AI RMF MAP 2.3 criterion — characterising data dependencies, data flows, and context of use for AI systems — applies equally in both contexts.

3.5 The Zero-Findings Majority and What It Does Not Mean

Thirty-two of the thirty-five institutions in the corpus produced zero findings. It does not mean these institutions have no AI governance gaps, and it does not mean their AI systems are well-governed. It means that no AI-related governance signals were observable from a public web surface scan of their primary URLs at the time of capture.

Several explanations are consistent with a zero-findings result: the chatbot co-locates disclosure and policy links (meeting the engine's threshold); the chatbot widget loads via JavaScript after the page capture window (making it invisible to the static DOM snapshot); the institution has configured its chatbot vendor's widget to include a privacy notice; or the institution has no observable AI capability at the public surface at all.

For institutions with confirmed public chatbot deployments — such as Del-One FCU (Breezy), WEOKIE FCU (Okie), TruStone Financial (Ruth), FirstLight FCU (Luna), and others — a zero-findings result suggests either that their chatbot surfaces include the necessary disclosure elements, or that those elements load dynamically and were not present in the static page capture. Both outcomes are governance-positive relative to the Harvard FCU finding, but neither constitutes a clean-bill determination.

4. Regulatory Signal Mapping

Signal-oriented crosswalk only. This section maps confirmed surface findings to the governance frameworks most likely to require follow-up investigation. It is not a compliance certification, legal attestation, or regulatory determination. No scoring is applied.

Finding Category	Framework / Lens	Governance Signal
Client-side Application Identifier Review	Governance Baseline	The purpose of the visible identifier cannot be determined from surface evidence alone. Institutions should document the owning SDK or service and confirm whether the endpoint is public, internal-only, or gateway-protected.
Agent Backend Exposure	Governance Baseline	Not triggered in this corpus. Backend mediation for any observable agent surface cannot be evidenced from surface scan alone.
AI Disclosure Gap	GAAIA §111 Transparency	Visible AI interaction without matching disclosure evidence may weaken transparency demonstrability under proposed GAAIA disclosure requirements.
AI Disclosure Gap	NIST AI RMF GOVERN 4.2	Visible AI interaction without matching disclosure evidence may weaken proof that AI risks and impacts are communicated to affected parties.
Data Governance Gap	GAAIA §131 Data Requirements	Visible user data-entry without matching policy or processing evidence may weaken data governance demonstrability under proposed GAAIA data requirements.
Data Governance Gap	GLBA / CFPB Chatbot Guidance	Visible user data-entry through an AI-assisted interface without co-located disclosure is relevant to CFPB guidance on chatbot use in consumer finance (2023) and GLBA privacy notice obligations.
Data Governance Gap	ISO/IEC 42001 Clause 6.1	Visible user data-entry without matching policy evidence may weaken proof that data-entry risks are assessed and treated within the AI management system.
Data Governance Gap	NIST AI RMF MAP 2.3	Visible user data-entry without matching policy evidence may weaken proof that AI data flows and context of use are mapped and characterised.

5. Conclusions

1. The dominant observable AI surface is concentrated outside the largest institutions

None of the eight US G-SIBs produced any finding in this scan. Their AI systems — Bank of America's Erica, Wells Fargo's Fargo, Capital One's Eno (in its authenticated form), and others — operate behind authentication or within mobile applications inaccessible to a public-vantage surface scan. The observable AI surface in the US banking market is dominated by community banks and federally insured credit unions deploying third-party AI chatbot platforms (interface.ai, Posh AI, Eltropy) directly on their public-facing websites.

2. The only substantive governance signal was a Data Governance Gap at a credit union

Harvard Federal Credit Union produced the sole medium-severity finding: a Data Governance Gap triggered by a file upload data-entry surface without a co-located privacy or terms link. This finding is directly relevant to NIST AI RMF MAP 2.3, ISO/IEC 42001 Cl. 6.1, GAAIA §131, and GLBA / CFPB chatbot guidance. The finding does not constitute a regulatory violation; it constitutes a public-surface governance signal requiring follow-up.

3. Low-confidence identifier signals require documentation, not remediation urgency

Citadel Credit Union and Capital One each produced a Client-side Application Identifier Review finding — the lowest-severity category. In both cases the evidence is attributable to standard third-party telemetry SDKs: Azure Application Insights (Citadel) and New Relic NREUM (Capital One). These findings should prompt internal SDK documentation review, not governance remediation.

4. The absent finding categories reveal a structural disclosure gap

No AI Disclosure Gap and no Agent Backend Exposure findings were detected in the corpus. The absence of AI Disclosure Gap findings is notable given that 35 institutions were selected for having observable AI chatbot surfaces. This suggests that, at the time of scanning, the chatbot UIs either include sufficient disclosure elements or load them via JavaScript after the static page capture window — a methodologically important distinction. Institutions with confirmed chatbot deployments producing zero findings should not be assumed to be in a stronger disclosure posture than Harvard FCU; the difference may be one of render timing rather than governance substance.

5. The regulatory framework gap is itself a governance signal

The absence of an enacted federal AI-specific statute with deployer-level disclosure obligations means that US institutions operating public-facing AI chatbots face a different regulatory risk posture than their European counterparts under the EU AI Act. The GAAIA discussion draft, if enacted in its current form, would introduce transparency obligations directly comparable to EU AI Act Article 50. Institutions that prepare for GAAIA-equivalent obligations now — by co-

locating AI disclosures and data-processing notices with their public AI interfaces — will be better positioned for the transition to a mandatory regime.

6. Surface evidence is a starting point, not a conclusion

Every finding in this report is produced under conditions of high evidence fragility. Surface scans capture what is observable at a single point in time from a public vantage point. They cannot evidence backend mediation, internal governance controls, staff training, risk assessments, or policy documents that exist off-surface.

6. Recommended Actions

Client-side Application Identifier Review — Citadel CU, Capital One

- Identify the owning SDK or service for the visible applicationID field and document its purpose.
- Confirm whether any related endpoint is publicly accessible, internal-only, or gateway-protected.
- If the identifier is attributable to a monitoring or analytics SDK (e.g. New Relic NREUM, Azure Application Insights), document this explicitly and close the finding.
- If the identifier is attributable to an AI or agent service, escalate to the Agent Backend Exposure remediation track.

Data Governance Gap — Harvard FCU

- Attach data-processing approval, data owner, and permitted-use evidence for user inputs submitted through the AI-assisted interface, including file uploads.
- Ensure a privacy policy or data-processing notice link is co-located with any data-entry surface associated with AI functionality, visible without requiring the user to navigate away.
- Review the finding with the responsible data protection or privacy officer and confirm alignment with GLBA privacy notice obligations and applicable CFPB chatbot guidance.
- Confirm whether the Posh AI platform terms include user-facing data-processing disclosure and, if so, ensure that disclosure is surfaced at the point of data entry.

All Institutions with Public AI Surfaces

- Confirm that any AI disclosure obligation under CFPB chatbot guidance and applicable state AI transparency laws is met co-located with the AI interaction surface, not only in a general privacy policy linked from the footer.
- Assess whether JavaScript-rendered chatbot widgets include disclosure elements in the rendered DOM state; static page captures may not reflect the full disclosure posture of dynamic widgets.
- Monitor the GAAIA legislative process and assess readiness to meet proposed §111 transparency and §131 data-requirements provisions ahead of any enactment date.
- Retain individual AIGovXRay scan reports with their Continuity Capsules to enable drift analysis and comparative assessment in future scan cycles.

7. From Observable Signal to Audit-Grade Evidence

This report is a Web Surface Scan: it captures what is observable at the public page layer and produces governance signals. Every finding is classified as high fragility precisely because surface signals are the beginning of an investigation, not its conclusion.

The AIGovXRay workflow is designed to move institutions from observable signal to deterministic, audit-grade evidence — through a structured progression:

Web Surface Scan	Observable signals — what any regulator can independently verify from a public vantage point
▼	
Deep Scan	Runtime & backend evidence — deterministic, corroborated, audit-grade
▼	
Evidence Pack	Structured artefact set with hashes, confidence scores, and framework crosswalks
▼	
Continuity Capsule	Drift analysis across scan cycles — governance posture over time
▼	
Audit Ready	Exportable evidence set for regulators, auditors, and counterparties

A Web Surface Scan — such as this report — produces the first layer: observable signals that any regulator, auditor, or counterparty can independently verify from a public vantage point. It answers the question: 'What can be seen?'

A Deep Scan adds runtime and backend evidence, corroborating or refuting the surface signal with deterministic data. It answers the question: 'What is actually there?'

The Evidence Pack and Continuity Capsule transform findings into an exportable, audit-ready artefact set — structured for regulatory use, drift analysis, and counterparty review across scan cycles.

For institutions where this report has identified a governance signal — and for institutions where a zero-findings result may reflect render-timing rather than governance substance — a Deep Scan is the natural next step.

Contact Mechapiens to initiate a Deep Scan and begin building deterministic evidence.

8. Limitations and Scope Boundaries

- All findings are derived from Web Surface Scans only. No backend, runtime, server-log, or vendor-internal evidence was analysed.
- Each scan represents a single point-in-time capture. Page content, AI functionality, and governance signals may change after the scan date.
- The corpus was constructed on the basis of vendor-confirmed or institution-disclosed AI chatbot deployments on public-facing websites. It is not a statistically representative sample of the US banking market.
- The 32-entity zero-findings cohort may include institutions whose chatbot widgets load via JavaScript after the static page capture window, institutions whose chatbot surfaces include all required disclosure elements, and institutions with no AI surface observable at the scanned URL.
- Confidence scores cited in individual reports (0.60–0.62) reflect heuristic derivation from limited surface evidence. They are not statistical confidence intervals in the inferential sense.
- The Client-side Application Identifier Review finding is explicitly acknowledged as ambiguous. Institutions affected solely by this finding should not be characterised as having AI governance gaps without further investigation.
- No enacted US federal AI statute with direct deployer-level disclosure obligations was in force at the time of scanning. All regulatory signal mappings to the GAAIA are based on a discussion draft that had not been enacted.

About AIGovXRay

AIGovXRay is an AI governance scanning product developed by Mechapiens PC. It produces evidence-first assessments of AI system governance signals from observable surface evidence. For further information visit www.mechapiens.com.

Disclaimer

This report is provided for informational purposes only. It is not legal advice, regulatory guidance, or a compliance certification. Findings should be reviewed with qualified legal and technical counsel before any regulatory or operational action is taken.