AIGovXRay

Public AI Governance Signals Across European Banking

2026 — Revised Edition

Report scope: 102 Systemically Important Institutions (SI banks) and materially related entities across the European Economic Area, Switzerland, and the United Kingdom, scanned via AIGovXRay Web Surface Scan during Q2 2026.

Generated:
June 2026
Issued by:
Mechapiens PC
Product:
AIGovXRay Web Surface Scan

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

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

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. Conclusions are therefore limited to what is materially observable at the public web surface.

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 102 entities. The primary cohort consists of Systemically Important Institutions (SI banks) across the European Economic Area, with supplementary coverage of materially related holding companies, digital-banking subsidiaries, and asset-management arms where those entities operate customer-facing digital surfaces in scope of the EU AI Act or equivalent national transpositions.

Finding Categories

The scan engine produces findings in four categories relevant to this edition:

- **Client-side Application Identifier Review:** A client-visible application identifier was detected in frontend-delivered code, but the captured evidence does not by itself distinguish agent routing from benign SDK, monitoring, or telemetry usage. This is a lower-confidence signal requiring vendor or runtime context before any backend exposure inference is made. Severity: low.
- **Agent Backend Exposure:** A client-visible API endpoint, HTTP header, or identifier specifically associated with agent or chatbot proxy infrastructure was detected in frontend-delivered code, with enough contextual specificity (endpoint path, service name, or named proxy identifier) to distinguish it from generic SDK metadata. Severity: medium.
- **AI Disclosure Gap:** A visible AI interaction surface (chat UI, prompt input, or AI assistant widget) was detected without a co-located user-facing disclosure of AI nature or AI limitations on the same captured page. Severity: medium.
- **Data Governance Gap:** 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 on the same captured page. Severity: medium.

Governance signals (privacy policies, terms links, cookie notices, and similar) observed during scanning are noted as contextual evidence but are not themselves treated as findings.

Evidence Fragility

All findings carry a fragility classification of high. This reflects the surface-only evidence tier: 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 should be understood as a governance signal requiring follow-up investigation, not as a concluded determination of non-compliance.

Corpus Overview

Of the 102 entities scanned:

Cohort	Entities	Share
Entities with at least one finding	**12**	11.8 %
Entities with zero findings	**90**	88.2 %
Total corpus	**102**	100 %

The 90-entity majority with zero findings should not be read as a uniform governance clean-bill. Absence of findings under a web surface scan reflects one of two distinct states: either no AI-related surface artefacts were observable at the scanned URL, or 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.

Findings Detail: Institutions with Confirmed Signals

The twelve entities below produced at least one confirmed finding. Each entry is drawn from the corresponding individual AIGovXRay report. Three institutions (Česká spořitelna, National Bank of Greece, Piraeus Bank) carry a HIGH overall risk signal with findings of medium severity requiring immediate review. The remaining nine carry a LOW overall risk signal with a single low-severity finding and a recommendation for routine follow-up.

Institution	System URL	Risk	Finding Categories	Observable Signals
Česká spořitelna	csas.cz	HIGH	**Client-side App ID Review**, AI Disclosure Gap, Data Governance Gap	Chatbot code, AI chat UI + prompt input visible
National Bank of Greece	nbg.gr	HIGH	**Agent Backend Exposure**, AI Disclosure Gap	Named proxy endpoint + applicationId header; AI assistant UI visible
Piraeus Bank	piraeusbank.gr	HIGH	**AI Disclosure Gap**, Data Governance Gap	AI assistant UI + prompt input; no disclosure or policy link on page
Erste Asset Management	erste-am.com	LOW	Client-side Application Identifier Review	applicationID hidden input; purpose unclassified from surface alone
Erste Digital	erstedigital.com	LOW	Client-side Application Identifier Review	applicationID hidden input; purpose unclassified from surface alone
Erste Group	erstegroup.com	LOW	Client-side Application Identifier Review	applicationID hidden input; purpose unclassified from surface alone
Sparkasse Austria	sparkasse.at	LOW	Client-side Application Identifier Review	applicationID hidden input; purpose unclassified from surface alone

SLSP (Slovenská sporiteľňa)	slsp.sk	LOW	Client-side Application Identifier Review	applicationID hidden input; purpose unclassified from surface alone
Erste Bank Hungary	erstebank.hu	LOW	Client-side Application Identifier Review	Chatbot code (unconfirmed); applicationID hidden input field
BCR (Banca Comercială Română)	bcr.ro	LOW	Client-side Application Identifier Review	Chatbot code (unconfirmed); applicationID hidden input field
Nordea	nordea.com	LOW	Client-side Application Identifier Review	applicationID in New Relic NREUM telemetry config; no AI surface confirmed
Mambu	mambu.com	LOW	Client-side Application Identifier Review	applicationID in New Relic NREUM telemetry config; no AI surface confirmed

Finding Frequency Across the Corpus

Finding Category	Institutions Affected	% of Corpus	Notes
Client-side Application Identifier Review	**9**	8.8 %	Unclassified applicationID or SDK header; may be benign telemetry
Agent Backend Exposure	**1**	1.0 %	Named AI proxy endpoint + service identifier (NBG only)
AI Disclosure Gap	**3**	2.9 %	Visible AI UI without co-located disclosure (NBG, Piraeus, CSAS)
Data Governance Gap	**2**	2.0 %	Visible AI data-entry surface without policy link (Piraeus, CSAS)

Note: counts reflect distinct institutions, not finding instances. CSAS contributes to three categories. NBG contributes to two (Agent Backend Exposure and AI Disclosure Gap). Piraeus contributes to two (AI Disclosure Gap and Data Governance Gap). All other institutions contribute one finding each.

Notable Findings

Client-side Application Identifier Review — The Dominant Pattern

Nine of the twelve institutions with findings produced exclusively a Client-side Application Identifier Review signal. This is the lowest-severity finding category in the engine's taxonomy. The evidence in all nine cases is a generic applicationID field in frontend-delivered HTML or JavaScript, and the engine explicitly states in each report that the captured evidence does not by itself distinguish agent routing from benign SDK, monitoring, or telemetry usage.

Two of the nine — Nordea and Mambu — present a particularly clear illustration of this ambiguity: their applicationID arises from a New Relic NREUM browser monitoring configuration, a standard web analytics and performance SDK with no AI-specific function. The individual reports for both institutions assign a low confidence score (0.62, derived) and recommend identifying the owning SDK before treating the signal as backend exposure.

The remaining seven — Erste Asset Management, Erste Digital, Erste Group, Sparkasse Austria, SLSP, Erste Bank Hungary, and BCR — produce an applicationID pattern via a hidden HTML input field. This pattern is consistent with a shared frontend template or form-handling library across the Erste ecosystem, and again the engine assigns a confidence of 0.62 with an explicit caution against inferring backend exposure without vendor context.

> The Client-side Application Identifier Review finding signals that a hidden or embedded identifier exists in frontend code and warrants documentation. It does not signal that AI infrastructure is deployed, that an agent is exposed, or that a compliance gap exists. The recommended action in all nine individual reports is the same: identify the owning SDK or service and document whether any related endpoint is public, internal-only, or gateway-protected.

National Bank of Greece — Agent Backend Exposure

NBG is the only institution in the corpus to produce a confirmed Agent Backend Exposure finding — the higher-severity variant that goes beyond the generic identifier pattern. The scan identified a specific proxy endpoint path (/apis/Nbg.NetCore.AI.Agents.Proxy.Web/cc-lg), a named service identifier (Nbg.NetCore.AI.Agents.Proxy.Web), and an applicationId header embedded in frontend-delivered JavaScript. This combination provides sufficient contextual specificity to distinguish the signal from benign SDK metadata, and the individual report assigns a derived confidence of 0.74.

Alongside the Agent Backend Exposure finding, NBG also produced an AI Disclosure Gap: the scan detected a visible AI assistant UI (the chatbot named Sophia) without a co-located user-facing disclosure. Both findings carry medium severity and the individual report recommends immediate review.

The Agent Backend Exposure finding does not indicate that the proxy endpoint is unsecured or publicly exploitable. It indicates that the path, service name, and identifier are observable in frontend code from a public vantage point, and that backend mediation cannot be evidenced from the surface alone.

Piraeus Bank — AI Disclosure and Data Governance

Piraeus Bank's public-facing website displayed a prominently branded AI assistant ("Piraeus AI assistant") with a visible prompt input field, without a co-located disclosure of AI nature or an accessible policy link on the captured page. The scan produced two medium-severity findings: an AI Disclosure Gap and a Data Governance Gap. No applicationID or backend identifier was detected.

This configuration is directly in scope of EU AI Act Article 50 transparency obligations for deployers of AI systems that interact with natural persons. The AI interaction surface was confirmed by observable DOM evidence: the captured page contained a textarea input labelled chatbotInput with placeholder text in Greek, providing strong grounding for both findings.

Česká spořitelna — Three-Category Signal

CSAS produced the highest number of findings in the corpus: a Client-side Application Identifier Review (low severity), an AI Disclosure Gap (medium severity), and a Data Governance Gap (medium severity). The scan detected chatbot-related client-side code alongside a visible AI assistant chat UI and a prompt input, without a co-located AI disclosure or privacy link.

Unlike the NBG Agent Backend Exposure, the CSAS identifier finding is the generic low-confidence applicationID pattern — the individual report assigns confidence 0.62 and the same caution applies. The governance concern at CSAS is therefore concentrated in the two medium-severity findings: the visible AI interaction surface deployed without disclosure or policy accompaniment.

Regulatory Signal Mapping

Signal-oriented crosswalk only. This section is not a compliance certification, legal attestation, or regulatory determination. It maps confirmed surface findings to the governance frameworks most likely to require follow-up investigation. 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	Backend mediation for the observable agent surface cannot be evidenced from surface scan. Internal control and proxy mediation evidence required.
AI Disclosure Gap	EU AI Act — Article 50	Visible AI interaction without matching disclosure evidence may weaken transparency and user-information defensibility.
AI Disclosure Gap	ISO/IEC 42001 — Clause 8.2	Visible AI interaction without matching disclosure evidence may weaken proof that relevant AI use information is communicated.
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.
Data Governance Gap	EU AI Act — Article 10	Visible user data-entry without matching policy or processing evidence may weaken data governance defensibility.
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 AIMS.
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 are mapped.

Conclusions

1. The dominant surface finding is a low-confidence identifier signal, not confirmed AI deployment

Nine of twelve institutions with findings produced only a Client-side Application Identifier Review — a finding the engine itself treats with low confidence (0.62) and explicitly declines to classify as backend exposure without further vendor or runtime context. Two of these are attributable to a

standard third-party analytics SDK (New Relic NREUM). The other seven are consistent with a shared frontend template across one banking group. This finding category should be read as a prompt for internal documentation review, not as evidence of an AI governance failure.

2. Three institutions present substantive governance signals requiring immediate review

National Bank of Greece, Piraeus Bank, and Česká spořitelna each produced at least one medium-severity finding. NBG's Agent Backend Exposure is the only confirmed case of a named AI proxy surface observable from the public web. Piraeus and CSAS each deployed a visible AI interaction interface without the disclosure and policy accompaniment that EU AI Act Article 50 and aligned frameworks require. These three carry an overall HIGH risk signal in their individual reports.

3. AI transparency obligations are the most directly testable regulatory gap

Three institutions displayed a visible AI interaction surface without a co-located user disclosure. This configuration is the most directly testable gap relative to existing regulatory frameworks. A surface scan is sufficient to establish the presence of the gap; it is not sufficient to determine whether equivalent disclosure exists elsewhere in the user journey. The recommended response is to confirm the disclosure requirement, add visible disclosure co-located with the AI interface, and verify persistence across user journeys.

4. 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. Institutions with no surface findings may have deployed significant AI capability entirely within authenticated or backend environments. Institutions with surface findings may have robust controls not visible from the public web layer.

Recommended Actions by Finding Category

Client-side Application Identifier Review (9 institutions)

- 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), 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 below.

Agent Backend Exposure (1 institution — NBG)

- Review the exposed proxy endpoint path, service identifier, and header and confirm whether backend routing details should be hidden, renamed, or moved behind stronger mediation.
- Confirm whether the exposed endpoint is publicly accessible, internal-only, or gateway-protected.

- Provide backend access-control and proxy mediation evidence for the referenced agent service.
- Capture a stable DOM snapshot and screenshot alongside HTML for audit traceability.
- Provide native runtime traces to validate backend behaviour beyond the visible UI.

AI Disclosure Gap (3 institutions — NBG, Piraeus, CSAS)

- Confirm the user-facing AI disclosure requirement for the affected interface against applicable regulatory obligations (EU AI Act Article 50 or equivalent).
- Add or evidence a visible AI disclosure, limitation notice, or human-review expectation co-located with the AI interaction surface.
- Verify that the disclosure persists across all user journeys through the interface, not only on the landing page captured at scan time.

Data Governance Gap (2 institutions — Piraeus, CSAS)

- Attach data-processing approval, data owner, and permitted-use evidence for user inputs submitted through the AI interface.
- Ensure a privacy policy or data-processing notice link is co-located with any data-entry surface associated with AI functionality.
- Review the finding with the responsible data protection or privacy officer.

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 90-entity zero-findings cohort may include institutions that deploy AI entirely within authenticated surfaces, mobile applications, or backend systems not observable through a public-URL surface scan.
- This report is not a compliance certification, legal opinion, or regulatory determination. It is an evidence-based governance signal report intended to inform follow-up investigation.
- Confidence scores cited in individual reports (0.60–0.74) 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 by the engine as ambiguous. Institutions affected solely by this finding should not be characterised as having AI governance gaps without further investigation.

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.

