10 questions to ask every vendor — and the red flags to listen for.

Use this sheet in your next GRC / IRM / decision-intelligence vendor scorecard. Each row is a capability a buyer should probe; the red-flag column is what weak vendors reach for; the green column shows what ERIO ships out of the box so you can baseline.

CAPABILITY	ASK THE VENDOR	RED FLAG	ERIO'S ANSWER
Signal ingestion	How many live external feeds does the platform ingest and normalise out of the box?	**Red flag:** Vague "integration framework" with no live, named feeds.	**ERIO's answer:** 18 named feeds (CVE/CISA · OFAC/OFSI/UN/FSF · Companies House · GLEIF · GDACS · PHE CDR · ReliefWeb) + custom.
Supplier-graph depth	Does the platform model **n-tier** suppliers (Tier 2–4+), not just tier-1?	**Red flag:** Tier-1 register only — everything deeper is a manual spreadsheet.	**ERIO's answer:** Live n-tier graph via Companies House + GLEIF + optional D&B / Sayari BYO licence.
Risk quantification	Does the platform produce P10/P50/P90 £ loss curves per material risk?	**Red flag:** Qualitative heat-maps or a separately-licensed quant add-on.	**ERIO's answer:** Deterministic Open FAIR Monte-Carlo as a first-class citizen; curves refresh as signal moves.
Replayable audit	Can every decision be reproduced byte-identically from inputs + FX + seeds?	**Red flag:** Immutable log of *what happened* — but the inputs that drove it aren't captured.	**ERIO's answer:** Replayable assurance pack — inputs + FX + FAIR seed chained (Patent GB2613313.2).
Regulator readiness	Is DORA / NIS2 / SOC 2 Type II / ISO 27001:2022 / NIST CSF 2.0 scored continuously?	**Red flag:** Policy library + attestation workflows that lag the live posture by weeks.	**ERIO's answer:** Continuous readiness scoring + gap list for each named standard.
Decision orchestration	Does the platform fire webhooks when risk thresholds are crossed?	**Red flag:** Dashboards you have to visit — no push, no orchestration.	**ERIO's answer:** Tenant-scoped webhooks (variant-aware) wire into ServiceNow / Jira / Slack / email.
AI governance	Does the platform label AI-assisted output and keep humans in the loop?	**Red flag:** Opaque LLM calls, no citation or audit of the generated text.	**ERIO's answer:** Every AI-assisted field labelled; EU AI Act limited-risk framing; outputs replayable.
Tenant isolation	Is customer data isolated per-tenant across read + write + ML model use?	**Red flag:** Shared training surface — "your data makes the model smarter".	**ERIO's answer:** Hard tenant isolation; customer data never used for provider model training.
Evidence pack	Can a security / procurement team self-serve a live evidence pack?	**Red flag:** PDFs emailed by the vendor's sales engineer on request — week-long turnaround.	**ERIO's answer:** One-click NDA-gated Trust Centre download: 6 legal PDFs + 9 live evidence files.
Time-to-value	How long until the first £ board-ready output?	**Red flag:** Months of configuration before the first useful screen renders.	**ERIO's answer:** Day-one demo tenant with 142 signals, 9 risks, FAIR curves and a 1-page board pack.