Falconry360™

WHITEPAPER

ENTERPRISE GRC / FALCONRYX

# AI-powered governance intelligence

How artificial intelligence is transforming governance, risk, and compliance operations in enterprise organizations — and how to adopt it with the oversight regulators expect.

Govern | Anticipate | Comply | Withstand | Assure

CONTENTS / FIVE PARTS, SIXTEEN SECTIONS

# What this paper covers


PART I · THE PROBLEM

01 The state of GRC today

02 The language bottleneck

PART II · THE TECHNOLOGY

03 Defining governance intelligence

04 The connected data model

05 Technique-to-task map

06 Use cases across five pillars

PART III · THE CONTROLS

07 Responsible AI — ISO/IEC 42001

08 Failure modes & mitigations

09 The regulatory landscape

PART IV · THE PATH

10 Implementation roadmap

11 A maturity model

12 Measuring the program

13 The buyer's checklist

PART V · THE SOLUTION

14 A working scenario

15 Why FalconryX

16 Delivered in Falconry360

EXECUTIVE SUMMARY

# GRC is language work — and AI reads at scale

Regulations, policies, control libraries, evidence, findings — nearly every artifact a risk function produces or consumes is unstructured text. That is the medium modern AI handles best. Applied with discipline, AI reads, maps, and prioritizes; accountable people decide.

### THE SHIFT

From periodic, document-driven review to continuous, evidence-led oversight of risk and compliance posture.

### THE CONSTRAINT

AI operating inside a control function must itself be controlled — ISO/IEC 42001-aligned, with human approval gates on every decision of record.

### THE PAYOFF

Skilled hours move from re-mapping overlapping frameworks to judgment: risk appetite, response, and assurance.

### THE PATH

Assist, then automate with approval, then orchestrate. Autonomy is earned phase by phase; accountability never moves.

# Fragmented tooling produces structural failures

01 **Duplicated assessment** — one encryption control tested four times for ISO/IEC 27001, NIST CSF, SAMA CSF, and PCI DSS, because nothing links obligations to a common control library.

02 **Assurance that decays** — control effectiveness is known at audit time, then degrades invisibly until the next cycle. Between cycles, the board governs on stale data.

03 **Regulatory velocity** — every horizon scan means reading circulars, consultations, and amendments line by line, then hand-tracing impact to policies and controls.

04 **Risk described, not compared** — registers of high/medium/low ratings give boards no basis to weigh a cyber exposure against a third-party concentration risk.

05 **Evidence assembled by hand** — screenshots, exports, and email trails gathered per audit, per framework: the most expensive way to prove the same fact repeatedly.

02 / THE LANGUAGE BOTTLENECK

# The bottleneck in GRC is not judgment. It is reading, mapping, and assembling — high-volume language work that consumes the hours judgment needs.

# Three properties, or it is not governance intelligence

**Governance intelligence** is AI applied to the records of governance — obligations, risks, controls, evidence, findings — on a connected data model, under human authority. A language model bolted onto a document store has the first ingredient only.

1

### Connected data model

Obligations, risks, controls, assets, evidence, and findings exist as linked records — not documents. Every AI output can point to the records it used.

2

### Models fit to the task

Extraction, semantic mapping, anomaly detection, scoring, and generation are different techniques with different failure modes — deployed deliberately, not as one general chatbot.

3

### Human authority by design

The system proposes; a named owner disposes. Approval gates, override logging, and audit trails are part of the architecture, not a policy overlay.

# The data model matters more than the model

AI can only cite what exists as a linked record. In Falconry360, every governance object is a node in one graph — so an answer to "are we covered?" traverses real relationships instead of searching documents.

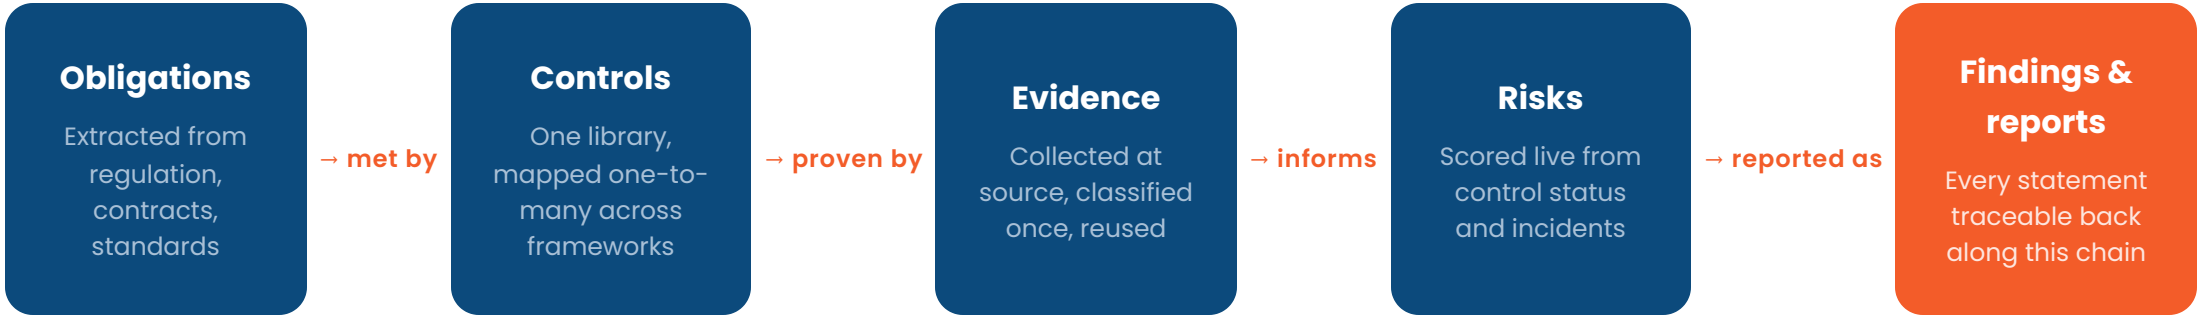


### Ask in either direction

"Which obligations does this control satisfy?" and "what breaks if it fails?" are the same traversal, reversed.

### Change propagates

A new obligation, a failed control, or expired evidence flags everything downstream of it — automatically.

### AI stays grounded

Generation is constrained to this graph. If a claim has no path through the records, it does not ship.

# Five techniques, five jobs, one human role

| AI technique | GRC task | Output | Human role |
|---|---|---|---|
| **Structured extraction (NLP / LLM)** | Obligation extraction from regulation, contracts, policies | Discrete, de-duplicated obligations with source citations | Approve or correct each extraction |
| **Semantic similarity (embeddings)** | Framework-to-control and obligation-to-policy mapping | Ranked candidate mappings with confidence levels | Accept, re-rank, or reject candidates |
| **Anomaly detection** | Continuous control monitoring over logs, configs, transactions | Drift alerts — disabled alerting, expired certificates, widening access | Triage, dispose, tune thresholds |
| **Predictive scoring** | Risk prioritization from control status, incidents, threat signals | Live exposure ranking with inspectable feature contributions | Set appetite, challenge the ranking |
| **Grounded generation** | Board reports, audit summaries, regulator responses | Drafts where every statement traces to a platform record | Edit, approve, and own the release |

# One AI layer, five connected pillars

| | |
|---|---|
| PILLAR 01<br>**Govern** | Policy drafting grounded in your own library · cross-policy conflict detection · attestation campaigns targeted to the roles a policy actually affects · delegations linked to the obligations they discharge. |
| PILLAR 02<br>**Anticipate** | Emerging-risk detection from incidents, near-misses, and key risk indicator drift — cross-read against consultations, threat intelligence, and supplier news · scenario analysis on demand, not once a year. |
| PILLAR 03<br>**Comply** | Obligation extraction on publication · one-to-many framework mapping against a single control library · gap analysis in hours, not quarters · Third-Party Risk Management (TPRM) questionnaires pre-answered from evidence and reviewed, not rewritten. |
| PILLAR 04<br>**Withstand** | Business-impact analyses kept current from live dependency data · recovery plans exercised against generated disruption scenarios · incident triage that classifies, routes, and drafts the first notification with regulator deadlines tracked. |
| PILLAR 05<br>**Assure** | Risk-based audit planning driven by live control and incident data · workpapers populated from the evidence repository · draft findings linked to supporting artifacts — auditors spend time on inquiry, not assembly. |

# Human oversight, structured by ISO/IEC 42001

| 42001 theme | In GRC operations |
|---|---|
| **Transparency** | Every mapping, score, and draft shows its sources and reasoning. No unexplained output enters the system of record. |
| **Human oversight** | 100% human approval gates on decisions of record — accepting a mapping, closing a finding, filing a report. |
| **Data governance** | Explicit rules on model data access, residency, and confidentiality — including where inference runs. |
| **Accountability** | Each AI capability has a named business owner. The model recommends; only the owner is accountable. |
| **Risk management** | AI systems sit in the risk register like any other material system — failure modes identified, controls assigned. |
| **Continual improvement** | Accuracy measured against human review outcomes; overrides feed evaluation. Quality is evidenced, not asserted. |

The practical test: when an auditor asks *"why did the system say this?"*, the answer — sources, model version, reviewer, approval — is retrievable in minutes.

# Four ways it goes wrong — and the control for each

### Fabricated mappings

A generative model asserts a control satisfies an obligation it does not — fluently and confidently.

**Control:** grounded generation only — outputs must cite platform records; uncited assertions are rejected at the gate.

### Automation bias

Reviewers approve machine suggestions by default; the approval gate becomes a rubber stamp.

**Control:** override rates monitored per reviewer; sampled blind re-review; confidence thresholds route low-certainty items to senior staff.

### Model & data drift

Regulatory language, business structure, or source systems change; yesterday's accuracy silently degrades.

**Control:** versioned models with evaluation baselines; accuracy tracked against human outcomes; re-validation on material change.

### Confidentiality leakage

Sensitive risk and audit data crosses residency or confidentiality boundaries through model inference.

**Control:** data-governance rules on model access and inference location; no training on client records without explicit agreement.

# Regulated in both directions

AI is simultaneously a tool that helps you comply and a technology you must comply about. A credible program addresses both from day one.

### GLOBAL

**EU AI Act** — risk-tiered obligations for providers and deployers, phasing in through 2027; AI used in regulated decisions needs classification.

**ISO/IEC 42001** — certifiable AI management system; the natural anchor for AI governance programs.

**NIST AI RMF** — Govern / Map / Measure / Manage, widely adopted as a program blueprint.

**DORA** — EU digital operational resilience for financial entities, including ICT third-party oversight.

### GCC

**SAMA** — Cyber Security and IT Governance Frameworks for Saudi financial institutions, with strong board-accountability expectations.

**NCA ECC** — Essential Cybersecurity Controls for national-scope entities.

**PDPL** — Saudi personal data protection; UAE and Qatar run parallel regimes, all with residency implications for AI processing.

**SDAIA guidance** — national AI ethics and adoption guidance shaping regulated-entity expectations.

One obligation mapped to many frameworks turns this landscape into a single connected library — **20+ frameworks pre-mapped** in Falconry360 — instead of parallel compliance programs.

# Adoption staged by autonomy, not technology

1

### Foundation & assist

MONTHS 0–3

Consolidate registers, controls, and evidence into one platform. AI in read-and-suggest mode; every output reviewed.

2

### Automate with approval

MONTHS 3–9

Extraction, mapping, and evidence classification as workflows with named approvers. AI governance stood up per ISO/IEC 42001.

3

### Continuous oversight

MONTHS 9–18

Source systems connected for continuous control monitoring and live risk scoring. Reporting becomes generated-and-approved.

4

### Orchestrate

MONTH 18+

A regulatory change triggers mapped control review, evidence refresh, and a board-briefing draft in one traceable thread. Humans still sign.

Each phase earns the next: measured accuracy against human review is the gate, not the calendar.

# Where does your program sit today?

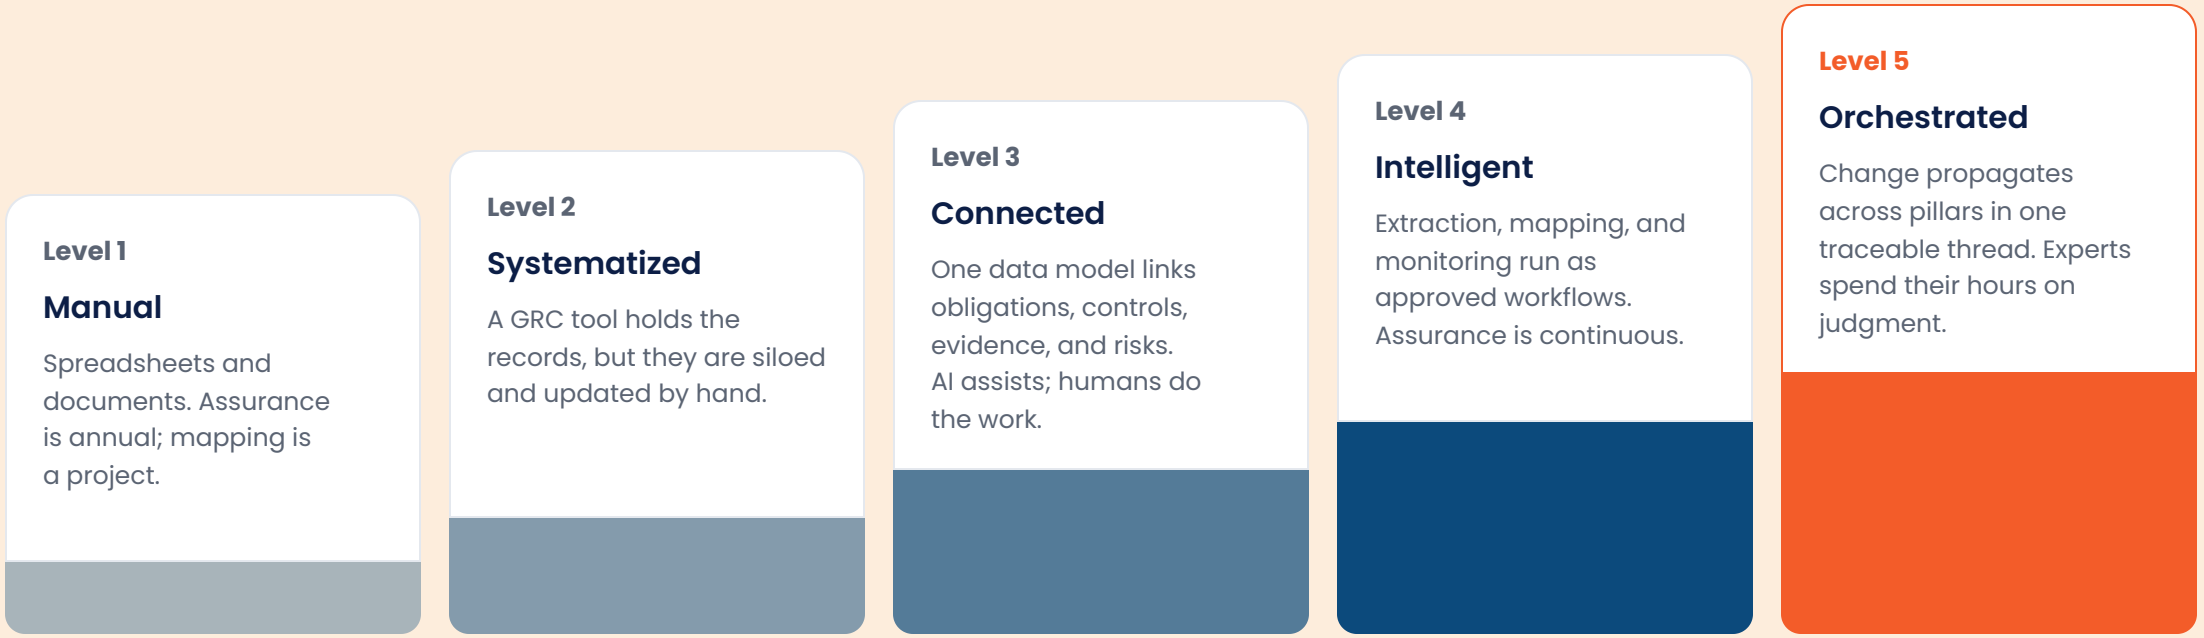


Most enterprises sit at level 1–2. The roadmap's four phases carry a program from wherever it starts to level 5 — with the approval gates widening only as measured accuracy earns it.

# Measure the machine like any other control

A governance-intelligence program reports on itself with the same discipline it brings to everything else. Six measures, all observable from the platform's own records:

### Time to map

Days from regulation publication to approved obligation-to-control mapping.

### Override rate

Share of AI suggestions corrected at review — tracked per capability and per reviewer.

### Evidence reuse

Frameworks satisfied per evidence artifact — the direct measure of collect-once, use-many.

### Coverage freshness

Age of the most recent effectiveness signal per control — the decay curve, made visible.

### Drift detection lead

Time between a control drifting and the program noticing — versus the audit cycle it replaces.

### Review concentration

Share of expert hours on judgment tasks versus assembly tasks — the ratio the program exists to move.

# Eight questions to ask any vendor — including us

1. Can every AI output show the records it used — sources, not summaries?
2. What is gated behind human approval, and can the gates be widened — or narrowed — per capability?
3. How is model accuracy measured, against what baseline, and who sees the number?
4. Where does inference run, where does data reside, and does that satisfy your residency obligations?
5. Is client data used to train models — and is "no" contractual or just current practice?
6. What happens to your mappings and audit trail if you leave — is the record exportable and legible?
7. Is the vendor's own AI governance aligned to ISO/IEC 42001 — and can they evidence it?
8. When the model is wrong, how do you find out — override tracking, drift alerts, or a customer complaint?

A vendor confident in its architecture will answer all eight in writing. Falconry360 does.

# An amended circular lands on Sunday morning

1

**Detect & extract**

Ingested on publication. 14 obligations extracted; 9 already covered by existing controls.

2

**Map impact**

5 new obligations mapped to policies, controls, and business units — with source paragraphs and confidence.

3

**Review**

Compliance lead reviews Monday, corrects one mapping, approves the rest. Corrections feed evaluation.

4

**Act**

Approved gaps become owned action items, deadlines aligned to the circular's effective date.

5

**Report**

Board-committee briefing drafted from the record; reviewed and issued by the CRO.

> The work compressed is reading, tracing, and drafting. The judgment stays exactly where it belongs.

# Not another AI feature. The intelligence layer governance runs on.

Point solutions bolt a language model onto a document store. FalconryX is different by architecture — and every difference on this slide is verifiable in a demonstration, not a claim to take on faith.

### Grounded in one graph

Every mapping, score, and narrative traverses the connected data model — obligations to controls to evidence to risks. If a claim has no path through the records, it does not ship.

### Native to all five pillars

One layer across Govern, Anticipate, Comply, Withstand, and Assure — so a regulatory change propagates to controls, evidence, resilience, and audit in a single traceable thread.

### Governed like a control

ISO/IEC 42001-aligned by design: 100% human approval gates on decisions of record, override tracking, versioned models, and an audit trail regulators can walk.

The result: the answer to "are we covered?" stops being a project and becomes a query — with the evidence attached.

# Delivered through FalconryX

FalconryX is not a separate product to integrate. It operates on the platform's connected data model — which is why its mappings, scores, and narratives stay traceable end to end.

**20+**

Frameworks pre-mapped to one control library

**100%**

Human approval gates on decisions of record

**5**

Connected pillars — Govern · Anticipate · Comply · Withstand · Assure

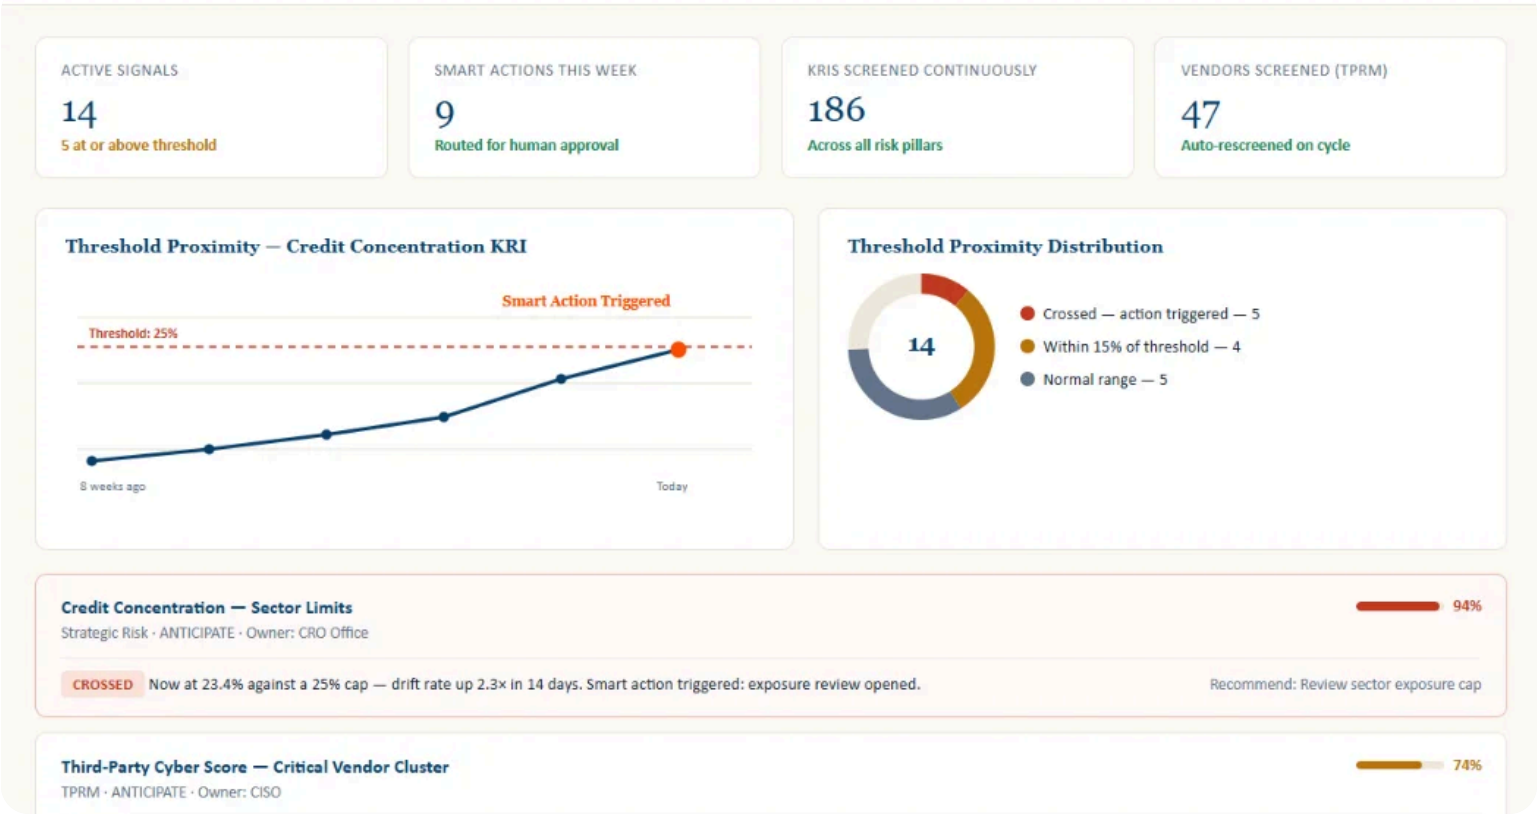

# From fragmented processes to connected execution

See how FalconryX applies governance intelligence to your frameworks, your risk register, and your evidence — with your oversight model intact.

**Request a demo**

contact@falconry.solutions · falconry360.com

Falconry360™