

Cognizant® Agent Governance for Insurance

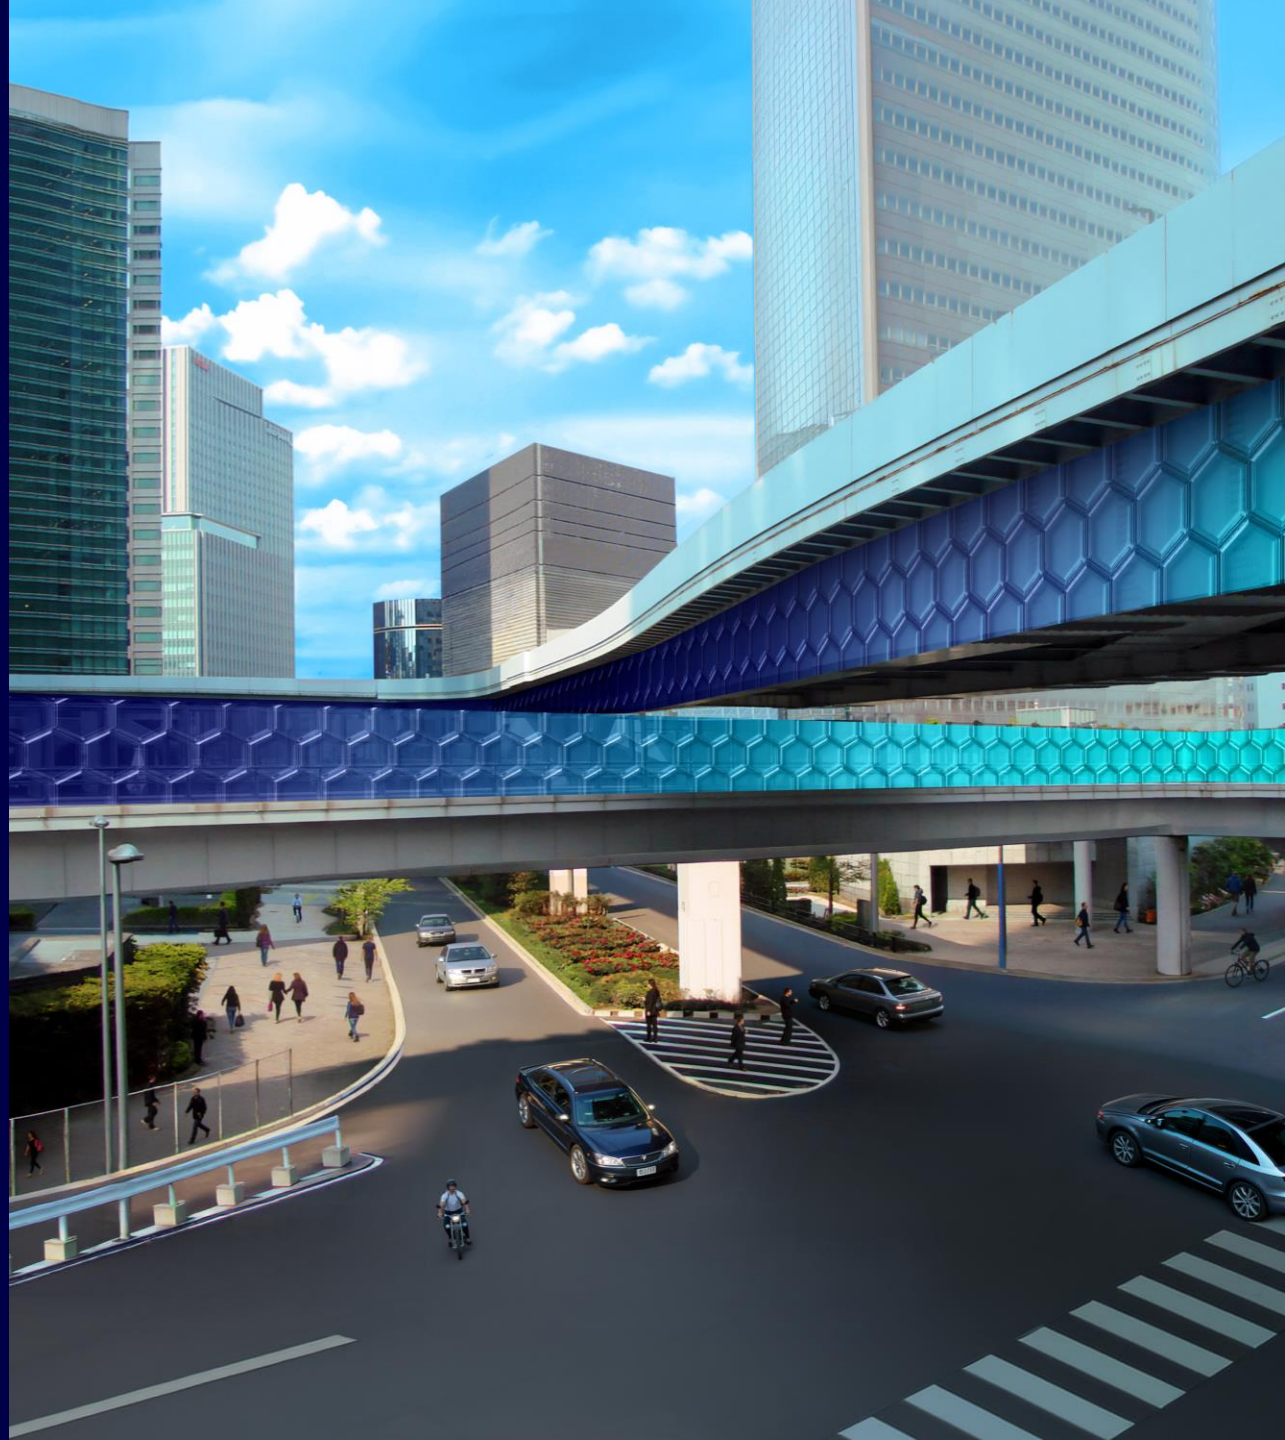
A unified, cloud-agnostic monitoring platform delivering real-time visibility into AI Agents and ML Models across any cloud ecosystem

Aug 2026

© 2025–2027 Cognizant | Private

Agenda

- AI Control Dashboard
- Four Strategic Monitoring Levers
- Business Need and Stakeholder Value
- Proposed Delivery Approach
- Business Benefits and Strategic Outcomes



© 2025–2027 Cognizant | Private

Business Need and Stakeholder Value

Move from fragmented monitoring to actionable insight by persona

Current gap

Organizations need a centralized mechanism to monitor model health, governance, compliance, and cost across a growing AI ecosystem.

Single-pane view
Proactive monitoring
Informed decision-making
Effective governance

IT Executives & Leadership

- Operational oversight
- Risk management
- Governance tracking
- Strategic decisions

Business Stakeholders

- Cost transparency
- ROI visibility
- Model effectiveness
- Outcome measurement

AI/ML Teams

- Performance monitoring
- Drift detection
- Reliability tracking
- Model optimization

© 2025–2027 Cognizant | Private

A Policy Stack the Enterprise Can Enforce

Separate principles, policies, standards, controls and evidence into one scalable model

Enterprise Standard

AI Governance Charter
Purpose, scope, risk appetite, non-negotiables and leadership commitments

Acceptable & Prohibited Use Policy
Approved, restricted, high-risk and prohibited AI uses

AI / GenAI / Agentic Risk Policy
Risk tiers by impact, autonomy, data sensitivity and decision criticality

Control Implementation

Lifecycle Governance Standard
Required governance gates from intake through retirement

Domain Control Standards
Data, model, prompt, oversight, security, privacy, vendor, agentic and audit

Evidence & Assurance Requirements
Minimum artifacts to approve, release, monitor, audit and recertify

Policies say what must be true | controls make it true | evidence proves it was true

© 2025–2027 Cognizant | Private

Governed Agent Lifecycle, from Design-Time through Re-certification

Re-certification triggers, Model version change, KB refresh, Data classification change, Capability change, Guardrail breach, Drift threshold, Security incident, regulatory rule update

INTAKE AND REGISTRATION

01 Intake
Use case, owner, LOB, users, outcome

02 Risk Tier
Agent, data, action risk, approval path

03 Register
Agent ID, framework, owner, version, identity

Artifacts: Use-case record, Risk tier, Agent registry entry, Service identity

BUILD-TIME GOVERNANCE

04 Adapt
Framework adapter to common contract, no functional rewrite

05 Classify
Data sensitivity, LOB and sub-tenant, purpose, retention

06 Govern Knowledge
KB approval, version, refresh trigger, source lineage

07 Train Gaps
Gap analysis, approved data, training report

08 Embed Controls
Prompts, guardrails, authority limits, Responsible AI controls

Artifacts: Adapter record, Data classification, KB version, Training report, Guardrail config

EVALUATION AND PROMOTION

09 Evaluate
Accuracy, Security, Infra, Regulatory Compliance, Data Governance, Guardrail, Capability, Auditability

10 Certify
Immutable certificate, capability list, training data link

11 Promote
Sandbox, Eval, Certified non-prod, Production

Artifacts: Evaluation scorecard, Certificate, Promotion approval Fail one dimension means no certificate

RUNTIME, FEEDBACK, RE-CERTIFICATION

12 Run & Enforce
Agent Core enforces certificate, policy, model, data, memory, capabilities, and audit logging

13 Monitor and Correct
Observability, challenger, HITL, incidents, drift / cost alerts

14 Re-certify or Publish
Re-test impacted agents, publish reusable skills and patterns

Artifacts: Mission trace, HITL record, immutable audit, decision lineage, feedback item, re-cert event, certified skill or pattern

AUTONOMY
Set by risk tier

L0 Assist
Answer, summarize, retrieve with evidence

L1 Recommend
Propose action with confidence and sources

L2 Act with approval
Prepare action, human approves execution

L3 Constrained autonomy
Execute within certified scope and limits

© 2025–2027 Cognizant | Private

Autonomy is a governance decision, determined by risk tier, data sensitivity, action type, capability declaration and required human approval

cognizant

Proving It Works: One Measurement Framework for the AI Estate

Is it running, is it right, and is it worth it: the three questions every AI product must answer, and the metrics that answer them

HEALTH (System, Operational, Behavioral)

Agent Runtime & Reliability
- Agent uptime (%)
- Agent crash rate / failure rate
- Average latency per step
- End-to-end workflow latency
- Timeout rate
- Error rate (system-level)
- Action/Tool failure rate

Model/Data Stability & Drift
- Model drift indicators (distribution shifts)
- Data drift (input feature distribution changes)
- Embedding drift (semantic space movement)
- Concept drift in business outputs
- Retrieval drift (declining retrieval precision/recall)
- Vector index freshness (recency of ingested content)

System & Integration Health
- API/Tool uptime
- Connector health (CRM, core systems, document APIs)
- RAG pipeline uptime
- Index latency (per KB)
- Cache hit/miss rate
- Memory store uptime (semantic memory, episodic memory)

Safety, Governance, and Compliance
- Guardrail violation rate
- Unsafe output flags (policy breaches)
- PII/PHI detection rate
- Jailbreak attempts blocked
- Proprietary info leakage detections
- Red-team stress-test results
- Output toxicity / offensive content scores

ACCURACY (Groundedness, Quality & Task Success)

Output Quality Scores
- Task success rate (per agent, per workflow)
- Output correctness (domain expert validation)
- Groundedness / hallucination rate
- Step-level correctness in multi-agent chains
- Incorrect response rate
- False positive/false negative rates (where relevant)
- Validation pass rate (structured or rule-based evaluation)

Retrieval & RAG Accuracy
- Retrieval precision & recall
- Relevance score (semantic similarity threshold)
- Over-retrieval / under-retrieval share
- Wrong document retrieval rate
- Context window utilization (%)
- Chunk attribution accuracy
- Citations correctness (hallucinated sources detection)

Multi-Agent Quality Signals
- Collaboration success rate
- Inter-agent contradiction rate
- Planning accuracy (correct decomposition)
- Delegation accuracy (task routed to right agent)
- Loop/no-progress detection frequency

Data & Label Quality
- Label correctness score
- Data quality defects (missing/incorrect fields)
- Content validation error rate
- Data freshness (how recent the source content is)
- Knowledge gaps detected by agents
- KPI alignment accuracy (outputs match business KPIs)

COST & EFFICIENCY (Token Use, Infra Costs, Efficiency)

Business Value Indicators (business-owned)
- Workflow cost per successful outcome
- Efficiency gain vs. baseline
- Throughput gain vs. baseline
- Operational cost reduction
- AI incident recurrence reduction (%)
- Business-outcome achievement rate (%)
- FTE hours returned to business
- SLA compliance rate

Effectiveness Metrics
- Tasks automated end-to-end (%)
- Manual intervention rate
- Average time saved per workflow
- Agent throughput (# tasks/hr.)
- Multi-agent throughput lift vs single agent
- Cycle time reduction (before/after agent)
- Human productivity increase (%)
- Consultant / analyst hours returned

Token Cost Metrics
- Total tokens consumed (prompt + completion)
- Cost per 1,000 tokens
- Cost per workflow / transaction
- “Cost of error” (tokens wasted on retries/rework)
- Token efficiency ratio (tokens per correct output)
- Prompt vs completion cost distribution
- Multi-agent overhead cost (% increase due to orchestration)

Infrastructure & Platform Costs
- Vector DB cost (storage + queries)
- GPU/compute hours consumed
- Orchestration framework cost
- Model hosting cost (per model, per version)
- API tool usage cost
- Embedding generation cost
- Memory store cost (long-term/episodic memory)

© 2025–2027 Cognizant | Private

Monitoring tiers ● Typically Monitored ● Need based Monitoring ● Advanced Monitoring

cognizant

AI Control Dashboard

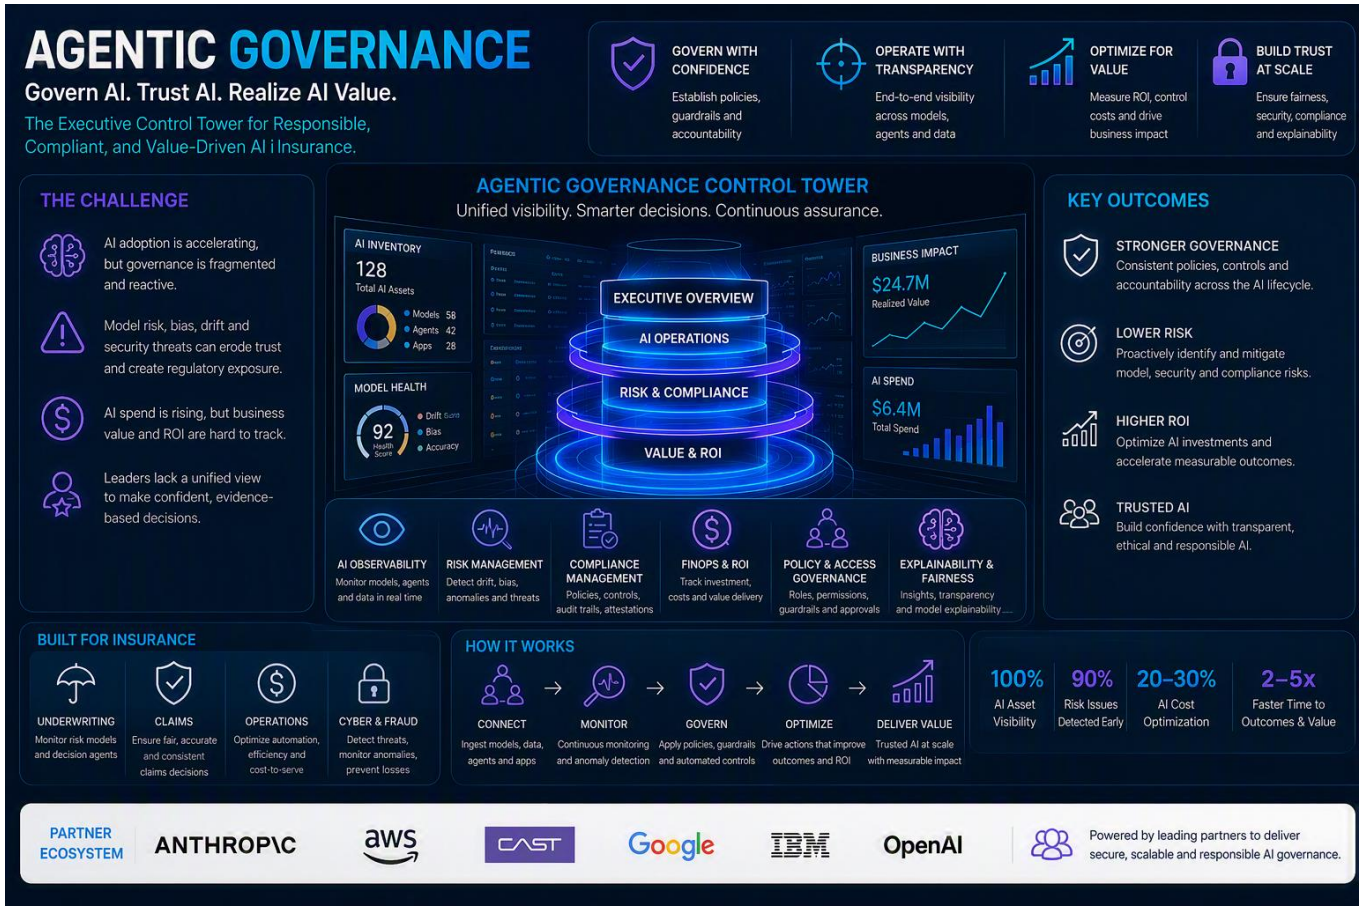
Unified model governance and observability across AI Agents, GenAI, and ML models

Why it matters:

Enterprises are scaling AI faster than their monitoring and governance mechanisms can handle, creating gaps across operational stability, model effectiveness, cost transparency, and compliance oversight.

Proposed answer:

A cloud-agnostic AI Control Dashboard delivers a single-pane view across AWS and Azure, giving leaders and teams real-time insight into model health, accuracy, cost, ROI, compliance, and governance.

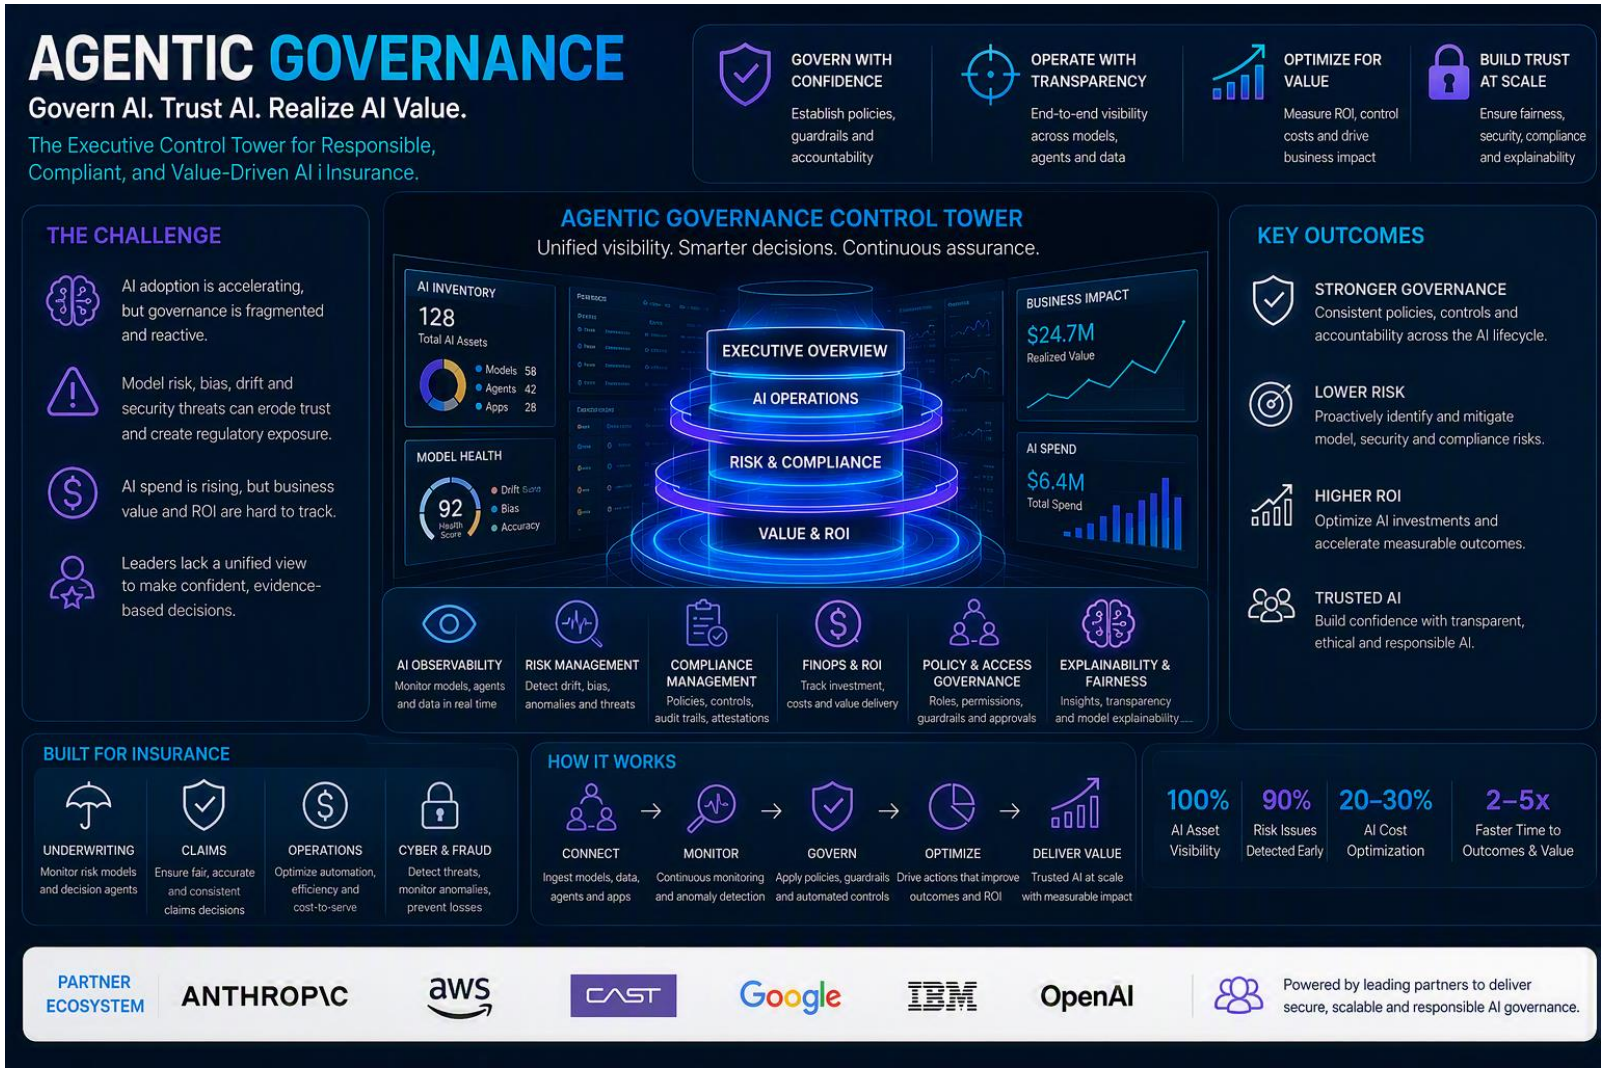


Cloud-agnostic | **Single-pane view** | **AI governance** | **Business value**

© 2025–2027 Cognizant | Private

AI Control Dashboard

Unified model governance and observability across AI Agents, GenAI, and ML models



© 2025–2027 Cognizant | Private

Proposed Delivery Approach

A pragmatic path from MVP foundation to enterprise operating model



MVP

Establish Foundation

Build core monitoring capabilities, onboard initial AI/ML models, and deliver foundational visibility across Health, Cost, Accuracy, and Compliance metrics.

Onboard + visibility



Scale

Grow Functionality

Expand model coverage, enhance observability, automate monitoring, and introduce advanced analytics and governance controls.

Automate + analytics



Institutionalize

Way of Life

Embed monitoring and governance into enterprise operations to drive continuous optimization, compliance, and data-driven decision-making at scale.

Operate + optimize

© 2025–2027 Cognizant | Private

Proposed Delivery Approach

A pragmatic path from MVP foundation to enterprise operating model

Business benefits

- ROI & business value tracking
- Operational risk reduction
- Real-time AI monitoring and visibility
- Faster leadership decision-making
- AI cost optimization and control
- Responsible AI and regulatory compliance
- Centralized governance and observability
- Improved trust and transparency

TEM lens

- Trust Gap: fragmented monitoring and governance
- Empathy Engineered: start with focused GenAI and ML/DL models
- Winning Moments: cloud-agnostic dashboard tailored by persona

PGR lens

- Protection: early detection of risk, drift and violations
- Growth: improved adoption, decisions and AI investments
- Retention: transparent visibility that strengthens trust and long-term value

© 2025–2027 Cognizant | Private



© 2025–2027 Cognizant | Private