cognizant | Collibra

Scaling AI with trust and control

Operationalizing AI oversight with Cognizant Trust™ Framework and Collibra AI Command Center

Executive summary

Organizations are accelerating the adoption of artificial intelligence (AI) to drive innovation, efficiency and competitive differentiation. Yet as AI systems scale—particularly with the rapid rise of autonomous agents that take action across enterprise systems—risk exposure expands just as quickly. Robust governance is essential to build stakeholder trust and to avoid lawsuits and regulatory penalties. The window to operationalize governance is narrowing as risks multiply: AI sprawl, shadow AI, ungoverned agents, opaque decision-making and emerging regulations all amplify exposure simultaneously.

"AI adoption is accelerating across enterprises, introducing new levels of scale, complexity and interdependence across models, applications and workflows." notes Gartner®.)1To unlock sustainable value while mitigating legal, regulatory and reputational risk, enterprises must move beyond ad hoc controls toward structured, lifecycle-based AI governance powered by real-time visibility and continuous trust signals. This requires clear accountability, embedded guardrails, automated traceability and end-to-end oversight across data, models and agents.

According to a Gartner research, "AI governance doesn't end at model deployment. Data leaders must take an ongoing approach to AI governance enforcement through automated feedback loops, early monitoring, and real-time controls to ensure continuous compliance and performance." [2] Cognizant Trust™ Framework, combined with Collibra AI Command Center, delivers a comprehensive approach that embeds governance into operational processes—transforming AI oversight from a reactive compliance function into a strategic enabler. Together, they help organizations move from AI ambition to execution with speed, confidence and integrity.

Business imperative for AI governance

As data becomes increasingly abundant and infrastructure more accessible, organizations are entering a new era of innovation—rapidly uncovering high-impact business opportunities and harnessing the transformative power of AI to reimagine what's possible.

The use of AI technologies in business processes and workflows presents immediate and multifaceted benefits. However, risks from unreliable or biased outputs that undermine trust—combined with latent vulnerabilities that expose organizations to security, ethical and regulatory issues—make AI adoption challenging. These challenges have intensified as AI itself has evolved. Predictive models brought concerns around bias and explainability. Generative AI introduced new risks: hallucinations, IP exposure and prompt injection. Now, agentic AI is widening the gap further. Autonomous agents no longer just generate content; they make decisions and trigger actions across enterprise systems. The result is a new class of AI control challenges that traditional governance models were never designed to address:

- **AI sprawl:** Models, agents and use cases proliferate across business units faster than central teams can track them.
- **Shadow AI:** Teams stand up AI initiatives outside of oversight, leaving gaps in inventory and accountability.
- **Opaque decision chains:** Agentic workflows span multiple systems, making it hard to trace what data, models and policies drove an outcome.
- **Continuous risk:** Static, annual review cycles cannot keep up with AI systems that update, retrain and act in real time.

AI deployed without effective governance presents immediate and far-reaching risks—from algorithmic bias and privacy violations to misinformation, security breaches and diminished human oversight. Left unchecked, these systems can inflict harm at scale and undermine public trust.

AI governance plays a critical role in the successful implementation and long-term adoption of AI initiatives. It is a critical enabler of reliable, ethical, transparent and compliant AI deployments across all industries. In highly regulated sectors such as banking, insurance and healthcare, AI governance ensures compliance and builds stakeholder trust. In industries like retail, manufacturing and energy, it serves as a catalyst for operational efficiency and innovation.

Implementing AI governance delivers dual business value by combining regulatory assurance with strategic growth. On one hand, it ensures compliance with evolving global frameworks such as the European Union Artificial Intelligence Act (EU AI Act) and Artificial Intelligence Risk Management Framework by National Institute of Standards and Technology, U.S. department of commerce (NIST AI RMF), mitigates legal and operational risks and reduces exposure to costly penalties. On the other hand, it acts as a catalyst for business transformation—driving operational efficiency, fostering stakeholder trust, enabling responsible innovation and creating sustainable competitive advantage.

Enabling AI governance at scale using Cognizant Trust™ Framework

Risks involved with AI technologies necessitate that organizations implement a companywide AI governance framework—including classifications, processes, risk categories and a central governing body for accountability across the lifecycle of every AI initiative.

Various regulatory bodies and private organizations have introduced their own governance frameworks—like the EU AI Act and AI Principles by OECD (Organization for Economic Co-operation and Development)—yet the underlying core principles remain largely the same. Cognizant Trust™ Framework builds on these core principles, integrating industry standards with our deep expertise in implementing governance structures and AI initiatives. The result is a proactive, comprehensive, repeatable framework designed to scale AI governance across the enterprise.

This framework provides:

- An internal and external governance structure that aligns and enhances existing organizational processes with ethical and regulatory standards—mitigating risks and ensuring compliance across the AI lifecycle
- Guidelines to build trust, ensure transparency, promote fairness, mitigate risks and empower users through safe and reliable solutions
- Tried-and-tested processes for scalable implementation of responsible AI practices across platforms and operations, incorporating robust compliance and risk management measures
- Coverage for the full spectrum of AI systems—from predictive models to generative AI to autonomous agents—so governance evolves with the technology, not behind it

For implementing governance for AI systems, organizations need to take a comprehensive approach that covers identification of the right governance framework, creating an inventory of their AI assets, and instituting a robust governance operating model. A process to identify evolving risks and continuously monitor AI initiatives, using well-defined metrics, must be an integral part of AI governance.

Cognizant helps clients operationalize AI governance through a framework-led approach, strengthened by our Collibra partnership, so organizations can scale AI quickly while maintaining safety and compliance.

Cognizant Trust™ Framework: Functional areas

AI governance functional area	Key activities
AI use case and business alignment	• Define and prioritize AI use cases based on business objectives, decision impact and expected value • Link each AI use case to its supporting assets (data sources, models, agents and workflows) • Track business KPIs and risk levels to guide governance, scaling and investment decisions across the AI lifecycle
AI governance framework assessment and setup	• Identify gaps in the current AI governance framework, processes and procedures • Develop a tailored governance framework and road map for managing AI- throughout its lifecycle, including predictive, generative and agentic system
AI asset inventory and mapping	• Register every AI model, agent and use case across business processes • Ensure lifecycle, compliance and performance metadata are recorded in a single source of truth to eliminate AI sprawl and shadow AI

AI governance functional area	Key activities
AI traceability and automated lineage	• Establish end-to-end traceability from AI use cases to models, agents, data sources and deployment environments • Automate lineage capture across AI assets to reduce manual documentation and operational overhead • Maintain continuous visibility into how data, models and agents are used within each AI use case • Enable impact analysis to quickly assess the effect of changes to data, models or configurations • Support auditability and accountability through consistent, system-generated traceability records
AI compliance and standards management	• Translate AI regulations and standards—including the EU AI Act, NIST AI Risk Management Framework and emerging AI UC-1 for autonomous AI agents (compliance framework built by Artificial Intelligence Underwriting Company)—into operational policies, controls and requirements • Assess and classify AI use cases based on regulatory risk, compliance obligations and required safeguards • Systematize compliance through standardized assessments, automated evidence collection and continuous compliance monitoring
Operationalizing the governance operating model	• Establish cross-functional governance committees and define roles and responsibilities across business units • Promote AI literacy and integrate governance into enterprise risk and quality systems • Automate policy enforcement, audit trails and enable seamless integration with existing AI and data workflows

Collibra AI Command Center as the AI control plane

Today, AI oversight is not just a compliance checkbox—it's a competitive advantage and a strategic imperative.

Scaling AI safely requires two layers working together—strong AI governance foundations at the asset level and continuous AI oversight across the enterprise. The foundation ensures every AI system—use case, model or agent—is structured, documented, classified and owned.

Continuous oversight ensures those same systems are monitored in real time as they run, change and act in production. Without the foundation, oversight has nothing to anchor to. Without continuous oversight, foundations age out the moment AI systems hit production.

Collibra AI Command Center delivers both. It is the end-to-end control plane for AI—bringing every AI agent, model and use case into a single governed view with continuous trust signals, automated traceability and proactive risk intervention—so enterprises can scale AI safely and prove its impact.

From governance foundations to continuous oversight

1. **Structure governance to fit your organization—the foundation, at the asset level**
   AI Command Center reflects existing practices, organizational maturity and AI risk levels. Collibra delivers code-first AI governance—capturing use cases directly from code and syncing them to a centralized registry—with flexible governance patterns and out-of-the-box assessments aligned to the EU AI Act, NIST AI RMF and AI UC-1. This establishes a trusted record for every AI system before it goes live.
2. **Operate AI systems with signals and controls—continuous trust and traceability**
   AI Command Center automatically traces AI execution and embeds governance controls across the AI lifecycle. It delivers automated cross-platform AI traceability across AI/ML environments, combined with a universal AI trust score that continuously quantifies readiness, risk and compliance in a single, actionable signal per system. This converts static documentation into a living view of AI as it evolves.
3. **Oversee AI continuously across the enterprise—continuous, at portfolio level**
   AI Command Center detects risks early and helps drive AI systems safely at scale. It provides a live portfolio-level view across every AI use case, model and agent—with interactive dashboards that translate governance signals into defensible decisions for boards and regulators.

Capabilities and supporting features

AI Command Center delivers the foundation and continuous oversight through six core capabilities, each grounded in the underlying activities that build trust, ensure compliance and guarantee reliability for AI initiatives.

1. One unified registry for every AI asset (foundation)

The Unified AI registry gives enterprises a single governed record for every AI use case, model and agent—with ownership, lifecycle stage, risk classification and business context attached. This single source of truth eliminates AI sprawl and shadow AI by ensuring every system is captured, documented and owned.

- Manage AI use case intake across the full AI lifecycle, connecting use cases with the models and agents that bring them to life
- Capture business use cases from across the organization—documenting problem statements, data requirements, expected outcomes and stakeholders
- Assign specific roles and responsibilities with clear ownership across every AI asset to foster collaboration and streamline processes
- Enrich AI asset context with detailed information on lifecycle process, data sources and operational metrics

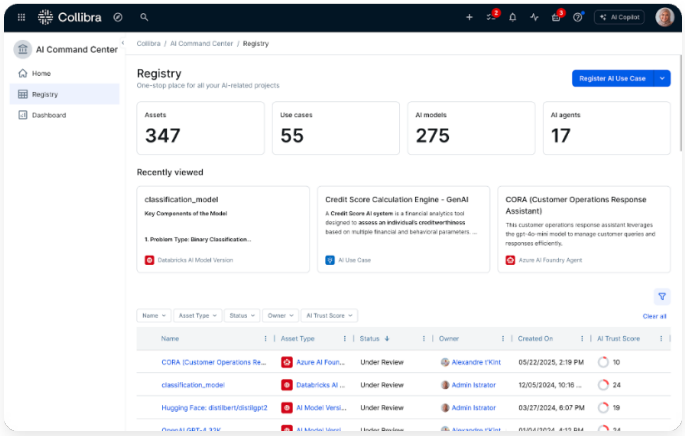

2. Code-first AI registration via CLI (foundation)

ML engineers register AI systems directly from the command line, embedding governance into the build workflow—not imposing it on top. Use cases sync to the centralized registry automatically, so governance is a by-product of building, not a separate exercise. Code-first AI governance is unique to Collibra and ensures the registry stays accurate as AI is built and updated.

3. AI UC-1 assessment templates, ready to run (foundation + continuous compliance)

Out-of-the-box assessment templates aligned to AI UC-1, the EU AI Act, and NIST AI RMF—launching May 2026. Prebuilt and ready to deploy without custom configuration, so teams evaluate AI readiness and compliance immediately, without translating regulatory standards independently.

- Drive domain-specific assessments across risk, data, legal and business dimensions using out-of-the-box templates
- Enforce governance controls by integrating policy checks into operational processes for continuous compliance
- Classify AI use cases by reliability and regulatory alignment to organizational and external standards
- Proactively manage legal and other risks by documenting known risks across data, models and agents and assigning risk levels to prioritize and manage them effectively

4. End-to-end AI traceability from data to agent (continuous)

Automated lineage across Databricks, Google Vertex AI, Amazon SageMaker, Azure AI Foundry and Snowflake—maintained without manual stitching. Every AI system is linked to the data it consumes and the policies it must comply with, making every output traceable, explainable and audit-ready.

- Connect to AI environments by integrating with ML platforms and pipelines to automatically collect metadata, lineage and operational metrics across AWS Bedrock, Amazon SageMaker, Google Vertex AI, Azure ML, Azure AI Foundry, Databricks, SAP AI Core and MLflow
- Develop comprehensive AI lineage across internal models, third-party models and agents
- Visualize AI traceability end-to-end across data, models and agents to understand how inputs flow, transform and drive AI outcomes

5. AI trust score: One signal per system (continuous)

The AI trust score aggregates documentation completeness, data quality, policy alignment and risk exposure into one continuously updated readiness signal per system. Teams know exactly when a system's readiness is changing and why—deployment and escalation decisions are based on live evidence, not assumptions or periodic reviews.

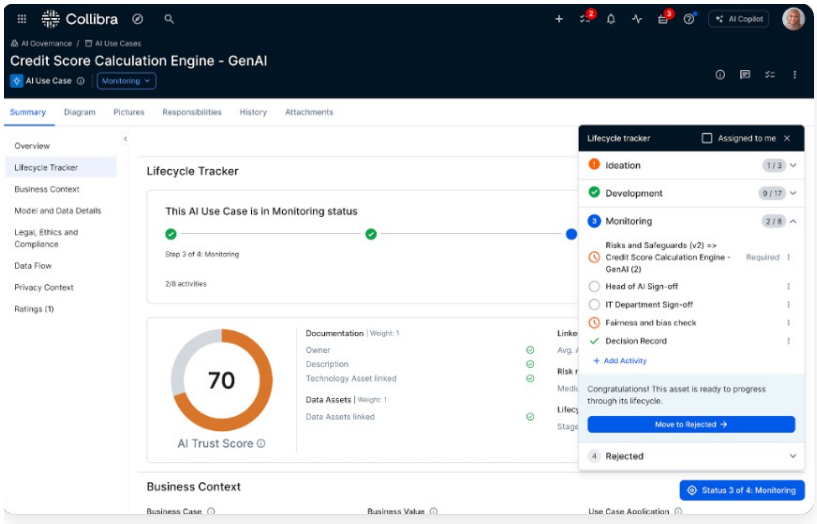

6. Interactive AI governance dashboards for leadership (continuous, portfolio level)

Interactive dashboards give leaders a 360° portfolio view across every AI use case, model and agent—organized by lifecycle stage, trust score and risk concentration. Leaders drill into exposure, prioritize investment and demonstrate AI accountability to boards and regulators based on live governance signals, not periodic reports.

- Oversee AI at enterprise scale through a unified registry of live AI systems (use cases, models, agents) with ownership, risk levels, approvals and exceptions—plus a single AI trust score per system, so leaders can approve, stage, remediate or retire AI systems with confidence.

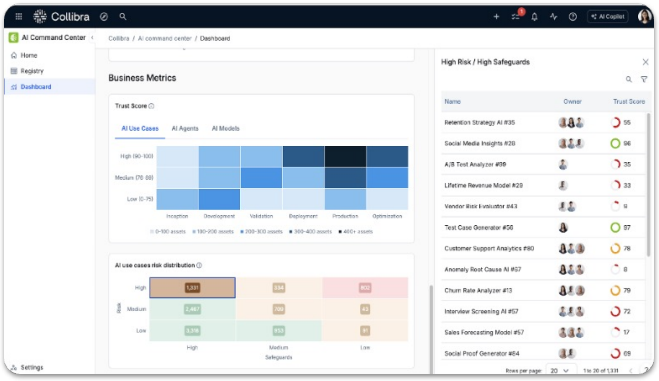

Conclusion

For organizations, the challenge with AI is not innovation—it's execution. Without clear ownership, trusted data and end-to-end oversight, promising AI initiatives stall before reaching production. The shift to agentic AI raises the stakes further. When AI agents trigger real decisions and actions across enterprise systems, periodic reviews are no longer enough. Enterprises need continuous, real-time oversight.

Robust AI governance—powered by a dedicated control plane—provides the foundation needed to move faster with confidence. By clarifying roles and accountability, aligning AI use cases to business objectives, managing risk proactively and ensuring transparency across data, models and lineage, organizations can operationalize AI at scale.

Together, Cognizant and Collibra deliver the frameworks and capabilities required to turn AI ambition into production-ready outcomes. Beyond meeting regulatory expectations, effective AI oversight enables teams to scale trusted AI at business speed, reduce friction between stakeholders and accelerate time to value. As seen in successful client deployments, governing AI systems doesn't slow innovation—it unlocks sustainable growth and operational excellence.

About our partnership

Cognizant has been a Collibra partner for over 10 years. We have implemented 100+ Collibra projects on the Collibra Platform and Collibra Data Quality across various cloud data platforms, applications and business use cases—spanning industries from financial services to healthcare to manufacturing.

Cognizant currently has over 250 trained Collibra practitioners with 10 Collibra Rangers globally, and has recently launched certification drives that will result in at least 20 Collibra Data Quality certified professionals and several more solution architects by the end of H1 2026.

Trusted data and AI oversight are foundational to building AI systems that are reliable, safe, ethical and scalable. With Collibra AI Command Center and Cognizant's deep expertise in this space, we help organizations build sustainable, scalable and trustworthy AI environments—across predictive, generative and agentic AI.

Explore the full set of AI Command Center capabilities at https://www.collibra.com/products/ai-command-center.

Author bios

David Talaga

David Talaga is a data and AI expert with a diverse background in analytics, product marketing, and AI governance. His career highlights include senior roles at Microsoft, where he helped businesses apply technology effectively, and positions at Talend and Dataiku, where he focused on data quality and operationalizing AI pipelines. Now at Collibra, he specializes in AI governance, helping organizations adopt AI responsibly by turning trusted data into enterprise-wide impact.

Dinesh Kumar Singh

Dinesh Kumar Singh is a Senior Manager in Data Governance practice at Cognizant and a recognized Data and AI Governance subject matter expert with over two decades of experience advising global organizations across multiple industries. He is passionate about advancing enterprise-scale governance frameworks that foster trusted data, responsible AI, and sustainable digital transformation.

Siddhartha Kanoria

Siddhartha Kanoria is a seasoned professional in the Data Governance and Data Privacy space, specializing in designing resilient data architectures that fuel AI innovation while ensuring rigorous compliance. His expertise lies in harmonizing Data Governance with emerging AI Governance frameworks, ensuring that the data fueling autonomous systems is not only high-quality but also ethically sourced and transparently managed.

References

1. Litan, A. (2026). AI Governance Requires More Than Policies. Gartner, June 10, 2026. Available at https://www.gartner.com/en/articles/ai-governance-trism
2. Priya Sundararaman, et al (2025). Operationalize AI Governance With Ongoing Postdeployment Control. Gartner, October 21, 2025. Available at https://www.gartner.com/document-reader/document/6947366?ref=pubsite

World Headquarters
300 Frank W. Burr Blvd.
Suite 36, 6th Floor
Teaneck, NJ 07666 USA
Phone: +1 201 801 0233

European Headquarters
280 Bishopsgate
London
EC2M 4AG
England
Tel: +44 (01) 020 7297 7600

India Corporate Office
Siruseri-Software Technology Park of India (STPI)
SDB Block—Ground Floor North Wing
Plot No H4, SIPCOT IT Park
Chengalpattu District
Chennai 603103, Tamil Nadu
Tel: 1800 208 6999

APAC Headquarters
1 Fusionopolis Link,
Level 5 NEXUS@One-North,
North Tower Singapore 138542
Phone: + 65 6812 4000

WF4965100