Bigeye

AI Trust Assessment Results

Emerging Stage

You're building AI oversight, but can't fully trust your systems yet.

What is AI Trust?

Most traditional data tools assume a human is thoughtfully deciding what data to access. AI agents operate differently — they can access vast amounts of data very quickly, combine it in unexpected ways, and make decisions about how to use it, at machine speed.

Without AI Trust infrastructure, you can't answer fundamental questions about your AI data usage:

- What data is this agent actually using?
- What sensitive data does the agent have access to?
- What decisions are being made by the agent using this data?

AI Trust means having visibility, accountability and control over how AI agents access and use your enterprise data, and ensuring that data meets quality standards for automated decision-making.

Introducing the
AI Trust Maturity Scale

To build and deploy AI responsibly you need a foundation of trust in the data that powers them. The AI Trust Maturity Scale helps teams understand where they stand today, and what it will take to move forward with confidence.

This scale isn't about theoretical best practices. It's about the real, operational steps that determine whether your AI initiatives can succeed and how much risk you're taking on along the way.

Your score reflects your team's current position on that scale. From there, you can begin identifying the biggest gaps, mapping realistic next steps, and building the infrastructure to support AI you can trust.

The Five Stages of AI Trust Maturity

Here's how the five maturity stages break down, from most mature to least mature.

Operational

You've built scalable guardrails around AI data usage with comprehensive oversight frameworks. Agents' data access are monitored and controlled with systematic rigor. Data quality, sensitivity classifications, and certification status are tracked and enforced automatically. AI decisions are auditable, and safe to scale.

Ideal business outcomes: AI initiatives launch faster with continuous governance. Scaling decisions are data-driven rather than risk-averse.

Success indicators: AI agents are automatically blocked from accessing unauthorized datasets before decisions are made, with real-time policy enforcement that validates every data request against sensitivity classifications and usage policies. Audit trails capture not just what data was used, but what agents attempted to access, including blocked requests and policy violations. When data classification or access rules change, updates are enforced across all AI systems, ensuring no unauthorized access incidents happen and. Teams are confident that agents operate only within their approved scope.

Managed

You've implemented structured data classification, quality monitoring, and approval processes across much of your data. Teams understand which data is validated for AI use, and consistent controls are in place to guide agent behavior.

Ideal business outcomes: Security and legal teams approve AI projects with confidence. Data-related incidents are rare and quickly resolved.

Success indicators: New AI use cases can be evaluated systematically against established data governance frameworks, with clear approval workflows that security and legal teams trust. When agents attempt to access data outside their approved scope, controls are in place to flag or restrict the action before inappropriate usage occurs. Data quality issues that could impact AI decisions are identified and resolved within defined SLAs, and teams can quickly determine whether specific datasets meet the standards required for AI consumption across different use cases.

Emerging (YOU ARE HERE)

You've built some of the necessary foundations: creating inventories, defining data characteristics, and introducing light controls around quality and usage. You're gaining visibility into your AI data landscape, but oversight is still limited and coverage is incomplete.

Ideal business outcomes: AI pilots move to production with fewer data-related surprises because teams have visibility into which datasets have been characterized and validated. Basic data controls can be demonstrated to stakeholders when questions arise about AI data usage.

Success indicators: Teams can identify gaps in coverage before they become blockers. Data inventories provide enough foundation to make informed decisions about which datasets are appropriate for specific AI use cases, even if comprehensive automation isn't yet in place.

Aware

You've recognized the need for oversight, but efforts are informal or inconsistent. Data stakeholders may be flagging risks manually, and there's growing concern about how data is used in AI, but no shared system to manage it systematically.

Ideal business outcomes: AI initiatives proceed with cautious optimism rather than fear. Internal teams understand and can articulate data-related risks.

Success indicators: Data issues are identified and discussed, even if resolution is inconsistent. Cross-functional teams collaborate on AI data concerns.

Unaware

You don't have many, if any, structured processes in place for data preparation, validation, or approval for AI use. Your organization may not even be aware of what's missing or understand the right steps to take. This is essentially "stage 0" -- where every organization starts, with no structure around AI data management.

Ideal business outcomes: AI experiments begin, though scaling is blocked by unknown risks and stakeholder concerns.

Success indicators: The organization recognizes that data governance matters for AI success.

What your score means

You're in the Emerging stage of AI trust maturity.

At this stage, your organization has likely built some foundational processes for AI data management, though the specific mix may vary. You've probably moved beyond purely informal approaches to establish some form of systematic oversight - this might include governance committees, working groups, or regular cross-functional meetings. You're likely monitoring data quality in some capacity, whether through basic alerts, manual checks, or systematic tracking for key datasets. Your organization probably has some classification of sensitive data types, though coverage and consistency may vary across different systems and teams.

However, your processes typically aren't comprehensive or fully integrated yet. While you've made significant progress from informal, ad hoc approaches, you're still working to achieve complete coverage and systematic validation across all your AI data usage.

The real cost of staying here

While you've made substantial progress beyond informal processes, staying at the Emerging level creates specific risks as your AI usage grows:

Quality risk

You likely have some form of data quality monitoring, but gaps in coverage or systematic validation mean quality issues can still emerge in areas that aren't fully covered. Your monitoring may catch technical problems, but without comprehensive business fitness criteria, you might approve data that meets basic standards but isn't appropriate for specific AI use cases.

> **Example:** A product recommendation AI performed well during testing, but once live it used seasonal data that hadn't been flagged in validation. Customers were shown winter products in spring, causing conversion rates to fall by 9%. The misaligned recommendations led to an estimated $120,000 in lost sales before the issue was corrected.

Sensitivity risk

You probably classify some sensitive data types, but incomplete or inconsistent characterization means sensitive information can still slip through your controls. Whatever scanning or discovery tools you use may work well when applied, but inconsistent application across all AI data sources creates blind spots.

Example: An AI analytics tool pulled customer email addresses into a report because sensitivity labels weren't applied at the column level. The report was later shared with an external vendor, triggering a compliance review and delaying a planned partnership campaign by three weeks. The delay added roughly $40,000 in additional legal and operational costs.

Certification risk

You likely have some criteria for validating data for AI use, but inconsistent application or incomplete documentation can lead to approvals without full validation. Your data tracking processes, whatever form they take, may help with some oversight but might not provide complete visibility into what data sources are being used or when they change.

Example: A forecasting AI was approved with a dataset that still included discontinued products. The forecasts overstated revenue by nearly 8% for the quarter, requiring finance analysts to spend over 150 additional hours reconciling projections with actuals.

The fundamental challenge of the Emerging stage is that partial implementation creates false confidence. Your controls work for what they cover, but gaps in coverage can create blind spots that become more dangerous as AI usage scales.

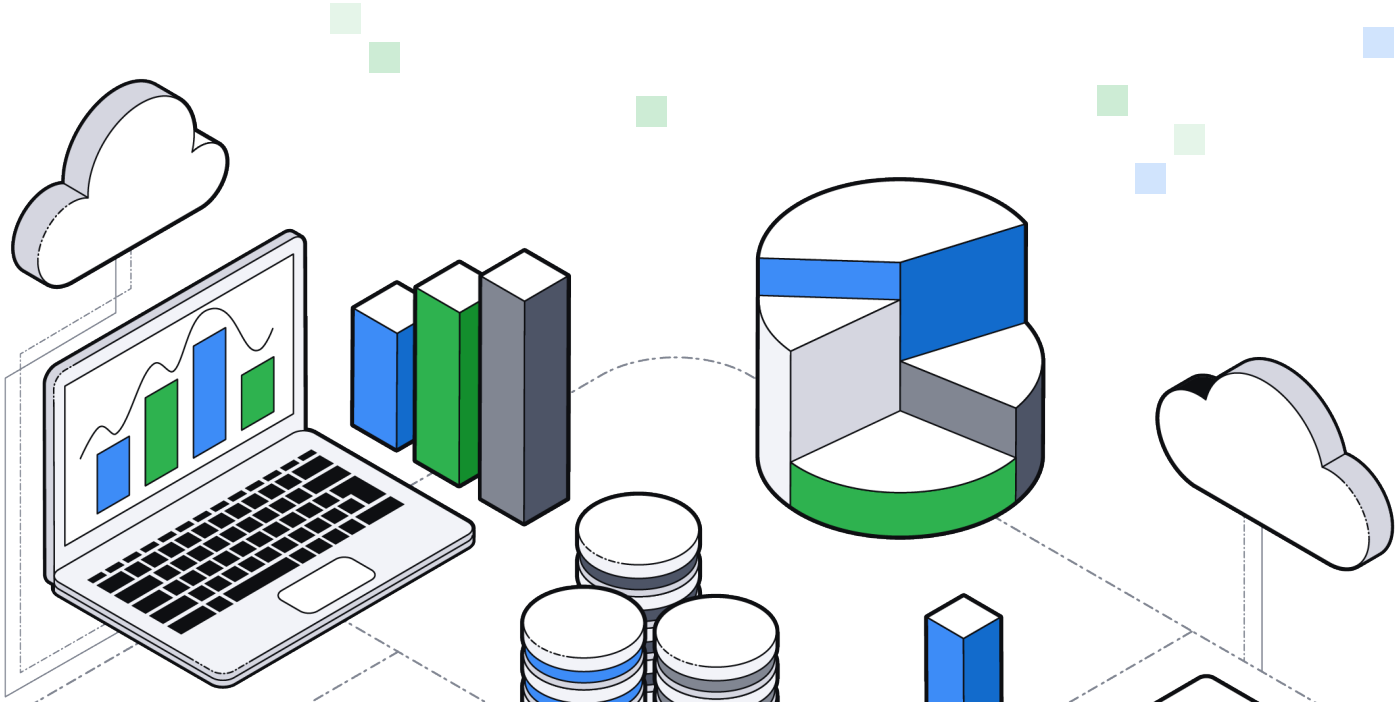

How the most mature orgs are operating

Organizations building AI trust into their infrastructure are moving beyond visibility. They're implementing enforceable controls that align to real business and compliance risk — without slowing down delivery.

Capability	Unaware	Aware	Emerging	Managed	Operational
Data Visibility	No or little visibility into AI data usage	Awareness of what data sources AI has access to	Systematic inventory of AI data sources	Comprehensive data cataloging with ownership	Real time monitoring of what data is being accessed by AI
Data Controls	No or little restrictions on AI data access	Informal guidelines for data use	Light controls and basic approval processes	Structured policies with consistent enforcement	Automated guardrails with real-time enforcement
Quality Management	Quality issues generally discovered after problems occur	Manual quality checks intermittently	Regular quality monitoring for key datasets	Systematic quality validation processes	Continuous quality assurance integrated into workflows
Risk Assessment	Unaware of full scope of data-related AI risks	Recognizes risks but no systematic approach	Basic risk identification and documentation	Structured risk management with regular reviews	Comprehensive risk monitoring with automated alerts

The Complete Journey to AI Trust Maturity

To reach full operational maturity, organizations need to build capabilities across three core areas. Here's how they connect:

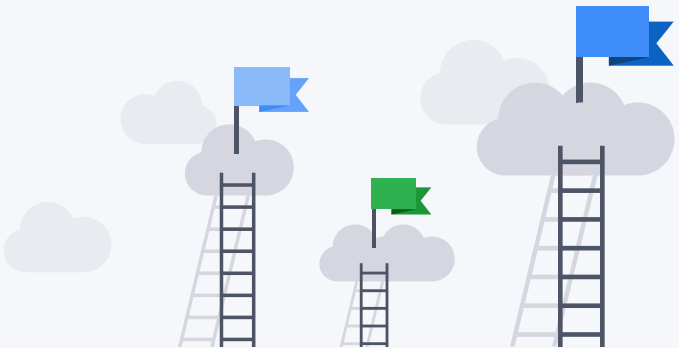

Quality Foundation: Know Your Data

Start by understanding what data your AI systems use and whether it's appropriate:

- **Map your AI data landscape** - Catalog which datasets, tables, and columns your AI systems access
- **Set usage boundaries** - Define which data fields should and shouldn't be available for AI consumption
- **Monitor data health** - Implement tracking for completeness, accuracy, timeliness, and consistency
- **Define fitness criteria** - Document what makes data appropriate for your specific AI use cases and create workflows to validate readiness
- **Handle problem data** - Build processes to identify and quarantine data that isn't ready for AI use

Sensitivity Controls: Protect What Matters

Control how sensitive information flows through AI systems:

- **Create approved inventories** - Maintain clear lists of datasets authorized for AI access
- **Classify for sharing** - Categorize all AI data by sharing permissions (Public, Internal, Confidential, Restricted)
- **Tag sensitive content** - Identify and mark PII, PCI, PHI, and other sensitive data types at the column level
- **Track privacy risks** - Maintain ongoing logs of privacy concerns and establish policies for sensitive data handling (like requiring SSN hashing)

Certification Framework: Validate and Approve

Ensure data meets all requirements before AI systems use it:

- **Combine quality and privacy assessments** - Document whether datasets meet both technical and sensitivity standards
- **Formalize approval workflows** - Create structured processes to validate and authorize data for AI use
- **Establish oversight** - Implement management processes aligned with industry standards for quality, security, and risk management
- **Maintain governance alignment** - Ensure practices meet established data management frameworks

The key insight: These concepts build on each other as you get more mature in your approach to AI ready data. You can't properly classify data until you know what you have, and you can't certify data until you understand both its quality and sensitivity profile.

Your Immediate Focus: Moving from Emerging to Managed

To move from Emerging to Managed, focus on completing and systematizing the capabilities you've already started building:

Quality Foundation

Complete Business Fitness Framework

Define comprehensive business rules

Build on your existing quality processes by documenting specific business requirements for different AI use cases

Integrate quality assessments

Whatever quality monitoring you have, expand it into systematic processes that combine technical and business fitness validation

---

Sensitivity Controls

Achieve Full Coverage

Complete data classification

Extend your current classification work to establish comprehensive sharing categories across all AI data

Systematize privacy risk management

Build on your current issue tracking to create comprehensive privacy risk registries

---

Certification Framework

Integrate All Assessments

Combine quality and sensitivity evaluations

Integrate your separate oversight processes into unified assessments that determine overall data readiness

What you already have: As an Emerging organization, you probably have foundational processes in place - governance committees, basic monitoring, classification systems, and approval criteria. Now it's time to complete the coverage and integrate these separate processes into comprehensive, systematic frameworks.

Where to start

Here's how teams at your stage typically complete their systematic capabilities:

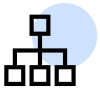

Extend your classification work

Build on whatever data classification you've already implemented. Whether you're using scanning tools, manual reviews, or systematic tagging, expand this to achieve comprehensive coverage across all datasets used by AI systems.

Complete your fitness framework

You likely have some quality monitoring or approval criteria in place. Extend this by defining clear business rules that determine when technically sound data is actually appropriate for specific AI applications.

Integrate your oversight processes

Whether you have governance committees, working groups, or other oversight mechanisms, work to integrate separate quality, sensitivity, and approval processes into unified data readiness assessments.

Systematize your issue tracking

Build on whatever issue reporting or tracking you currently have - whether that's ticketing systems, regular meetings, or documentation processes - to create comprehensive visibility into data-related problems and their resolution.

Document your governance alignment

Whatever governance structures you have in place, document how they align with established data management frameworks to create a foundation for more sophisticated compliance and risk management.

Why it matters now

Organizations that remain in the Emerging stage often face a specific challenge: their processes work well for what they cover, but incomplete coverage creates growing risks as AI usage expands. This typically creates two key problems:

Coverage gaps limit scaling - Your current processes may work effectively for key datasets or primary use cases, but gaps become more significant as AI usage grows across more parts of the organization.

Process fragmentation reduces effectiveness - You likely have separate processes for different aspects of data oversight, but without integration, this creates inefficiencies and inconsistencies that become more problematic as AI deployments become more complex.

The path forward involves completing the systematic coverage you've started and integrating your various oversight processes into comprehensive frameworks. Organizations that make this transition typically see more predictable approval processes, more effective risk management, and easier scaling as their comprehensive processes can handle increased complexity.

What to do next

Ready to complete your systematic frameworks?

Moving from Emerging to Managed requires completing the coverage of your existing processes and integrating your various oversight mechanisms into comprehensive assessment frameworks. Many teams find this easier with expert support.

[Get hands-on implementation support]

Our professional services team helps organizations build AI Trust capabilities from the ground up. We'll work alongside your team to create and implement the foundational controls you need to move confidently to the next maturity stage.

Not sure where your biggest blind spots are? [Talk to our team]

We'll help you assess your current state and map out a realistic implementation plan.

Want to understand the bigger picture first?

[Download the full AI Trust for Enterprise Organizations whitepaper]

Bigeye