Bigeye

AI Trust Assessment Results

Aware Stage

You've recognized the problem, but you're still managing it manually.

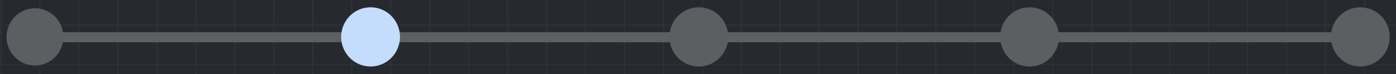

What is AI Trust?

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.

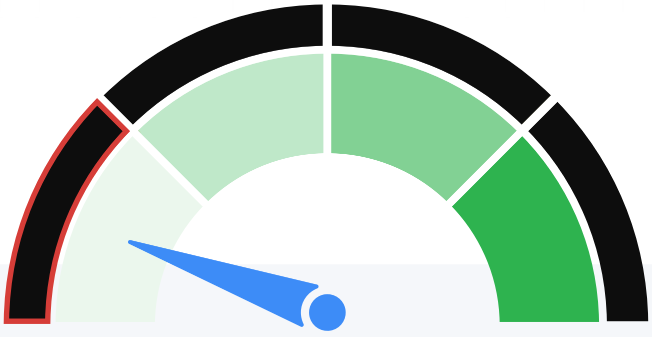

Introducing the AI Trust Maturity Scale

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'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 ARE HERE)

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 **Aware stage** of AI trust maturity.

At this stage, your organization has recognized the need for oversight around data used in AI, but your efforts are largely informal or inconsistent. You have some working groups and ad hoc processes, but they depend heavily on individual relationships and tribal knowledge. Teams may be flagging risks manually and having conversations about data quality, but there’s no shared system to manage it systematically.

You’re past the “we don’t know what we don’t know” phase, but you’re still relying on people rather than processes. And that creates its own set of risks.

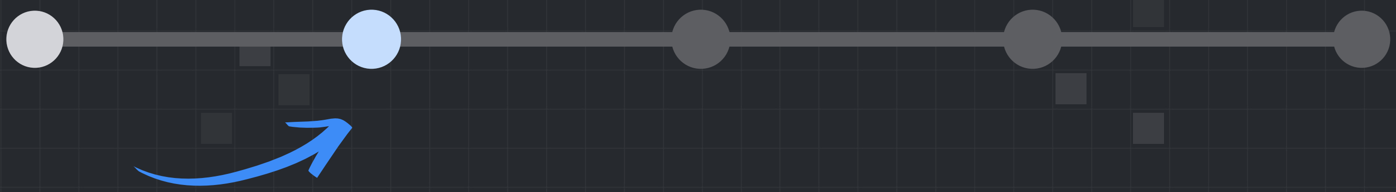

The real cost of staying here

While you’re more aware of risks than organizations in the Unaware stage, you’re still facing significant exposure because your controls depend on people, not processes:

Quality risk

Some teams manually review datasets before use, others rely on “tribal knowledge” about what’s reliable. Data quality issues get reported after they’re discovered, not prevented before they cause problems. One person’s departure can eliminate institutional knowledge about data reliability. Because reviews happen inconsistently, quality problems still slip through and get embedded in AI decisions.

Example: A marketing team’s AI model started underperforming after being trained on data from a deprecated source. With no process in place to validate dataset quality, the issue went unnoticed for two quarters. Campaign conversion rates dropped by 12%, resulting in roughly $180,000 in lost pipeline before the problem was traced back to the training data.

Sensitivity risk

Sensitive data is “labeled manually, but not consistently.” Some teams have their own policies, others rely on individual judgment. Without systematic classification, AI agents can still access data they shouldn’t. Privacy issues are “handled informally when discovered,” meaning violations may go undetected for extended periods before someone notices and reports them.

Example: An AI customer service tool accessed employee personal information stored in the same database as customer records. Because there was no consistent classification process, the system occasionally surfaced home addresses in interaction logs. It took several weeks to catch the issue, and the remediation effort cost the company over $50,000 in additional compliance work and legal review.

Certification risk

Teams choose datasets “based on past usage or personal judgment.” Training data documentation is “inconsistent” and kept in “informal notes.” When issues arise, ownership gets “assigned manually on a case-by-case basis.” This works until key people leave, priorities shift, or usage scales beyond what informal processes can handle.

Example: A finance AI began generating inaccurate revenue forecasts after pulling data from an A/B test that had ended months earlier. Because the dataset was never formally certified for production use, the issue persisted across three reporting cycles. The forecasting errors led to budgeting variances that required more than 200 hours of manual reconciliation, costing the business an estimated $75,000 in additional analyst time.

The fundamental problem is that manual processes don’t scale with AI adoption. As you deploy more agents and use cases, the informal relationships and ad hoc reviews that work today will become impossible to maintain.

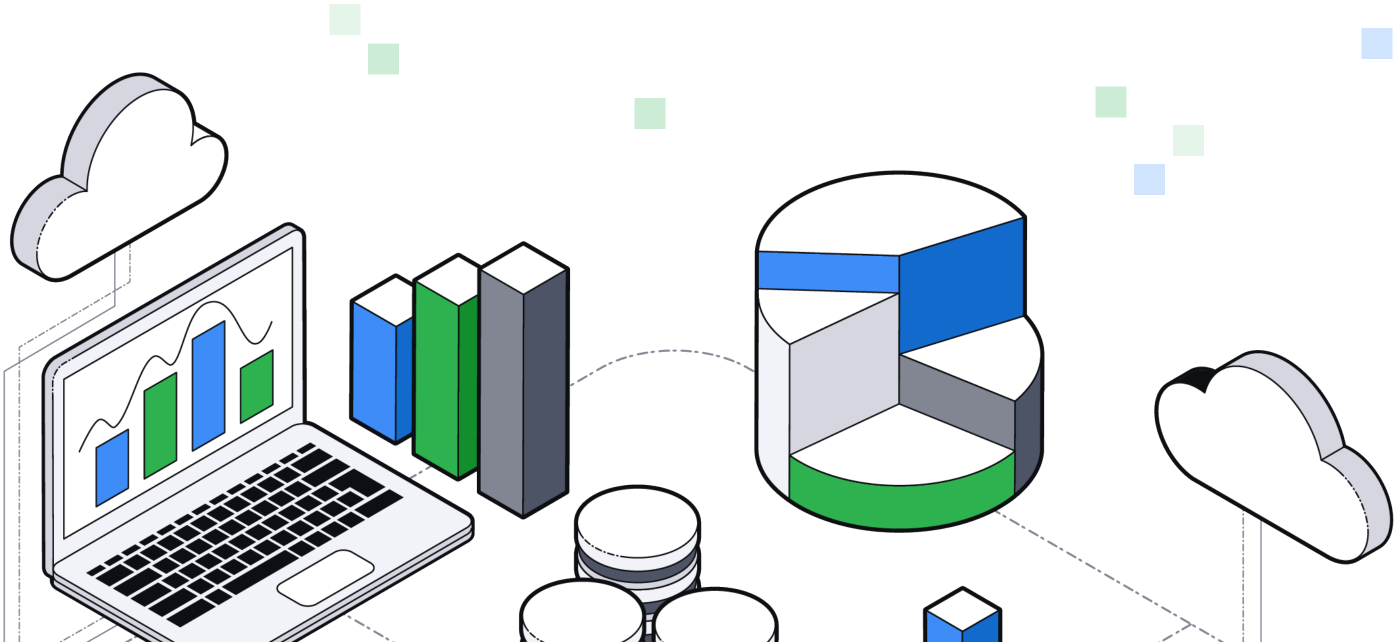

How the most mature orgs are operating

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:

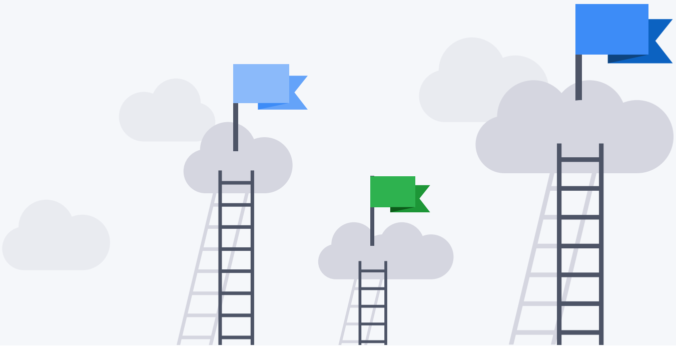

Quality Foundation: Know Your Data

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 Aware to Emerging

While the above represents the full scope of AI Trust maturity, you don't need to tackle everything at once. To move from Aware to Emerging, focus on building these specific capabilities:

Quality Foundation

Build Understanding and Processes

Establish data quality dimensions
Move beyond basic awareness to systematically understanding completeness, accuracy, timeliness, and consistency across your AI datasets

Create business processes for data readiness
Develop workflows that validate data meets your specific business requirements for AI use

Implement data triage processes
Build systematic processes to identify and quarantine data that isn't ready for AI use

Sensitivity Controls

Define Data Characteristics

Complete column-level data characteristics
Build on your existing data lists by defining PII, PCI, PHI, and PSI characteristics for each column your AI systems use

Certification Framework

Expand Governance Alignment

Align with data governance standards
Document how your AI data practices align with established frameworks like CDMC and DCAM, building on your existing process oversight

What you already have: As an Aware organization, you've established the foundational data inventories and basic approval documentation. You have process oversight in place. Now it's time to build systematic understanding and processes on top of that foundation.

Where to start

Here's how teams at your stage are building these emerging capabilities:

Turn data awareness into systematic understanding

You know what data you have and where it's used. Now establish consistent ways to measure and monitor its quality. Create standardized definitions for what "good enough" looks like across different data quality dimensions.

Build on your existing approvals

You're already documenting approval decisions. Extend this by creating systematic processes that define how data gets validated for business fitness and how unsuitable data gets flagged and handled.

Characterize your sensitive data systematically

Move beyond knowing which datasets contain sensitive information to systematically tagging PII, PCI, PHI, and PSI at the column level across all your AI data sources.

Formalize governance alignment

You have process oversight established. Now document how these processes align with industry data governance standards, creating a bridge between your current practices and comprehensive governance frameworks.

Why it matters now

Organizations that stay in the Aware stage often face a particular challenge: they know enough to worry, but not enough to act systematically. This creates two common problems:

Analysis paralysis - Teams know they need to establish data quality dimensions and characterize sensitive data, but get overwhelmed trying to implement comprehensive frameworks all at once. They have the data inventories but can't decide how to systematically assess and classify what they've catalogued.

False confidence - Having basic inventories and approval documentation feels substantial, but these manual processes break down as AI usage scales. Teams discover that knowing what data they have isn't the same as knowing whether it's appropriate for AI use when they try to expand beyond pilot projects.

The path forward isn't about building perfect systems overnight. It's about taking the awareness you already have and making it systematic enough to scale.

The organizations that build systematic processes early see practical benefits: security teams approve projects faster because they understand the data controls, launches have fewer surprises because data issues are caught earlier, and scaling becomes easier because the trust infrastructure can grow with AI adoption.

What to do next

Ready to systematize your informal processes?

Moving from Aware to Emerging requires turning tribal knowledge into documented processes, informal guidelines into systematic controls, and ad hoc reviews into regular workflows. 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