Bigeye

AI Trust Assessment Results

Unaware Stage

Your AI systems are running ahead of your trust infrastructure.

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.

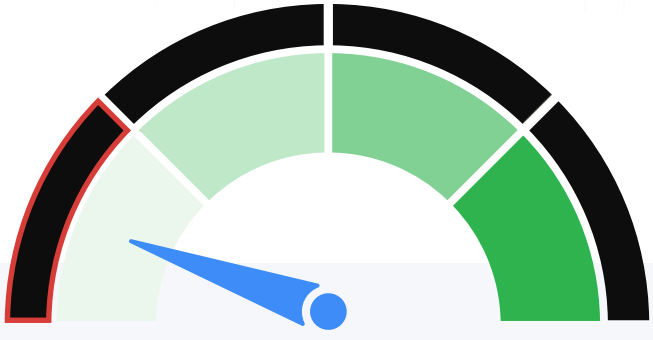

Introducing the
AI Trust Maturity Scale

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'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 ARE HERE)

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 **Unaware stage** of AI trust maturity, the first stage of your AI Trust journey.

At this stage, your organization likely hasn't formalized how data is prepared, governed, or approved for use in AI. There may be individual efforts happening across teams, but there's no shared system for tracking which data is safe, high-quality, or authorized for agent use, and no clear process for what happens when something goes wrong.

Without the right infrastructure, you're assuming risk you can't currently measure.

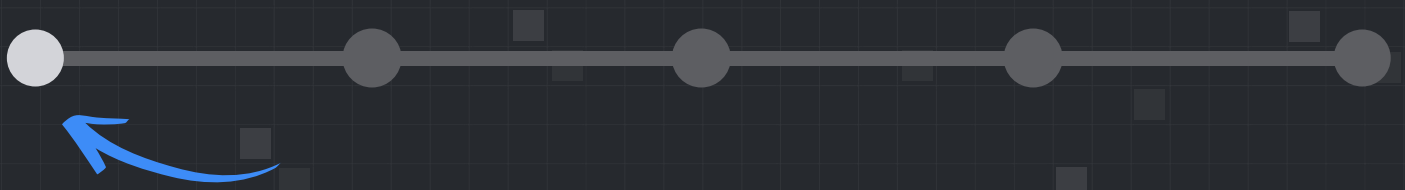

The real cost of staying here

AI systems make decisions automatically based on whatever data they can access, with no quality controls to catch errors before they impact business outcomes. Because everything happens "automagically," there's no chance to catch errors before they're embedded in decisions. Highly sensitive data can be processed and shared through AI without appropriate context-based restrictions, some data might be fine for internal use but inappropriate for broader sharing. Most critically, data that was never validated, intended, or certified for AI use gets pulled into models and decisions without anyone knowing.

If you're experimenting with agentic AI (or even just beginning to support basic AI/ML workflows), staying at this level of maturity carries three critical risks:

Quality risk

Agents can surface stale, partial, or inaccurate data with total confidence. There's no monitoring in place to flag what's out of date or out of scope. Because AI operates automatically at machine speed, errors get embedded in decisions before anyone can catch them. One data quality issue can derail a forecast, a workflow, or a customer interaction.

Example: A sales team's AI assistant consistently recommended outdated pricing from a dataset that hadn't been refreshed in a few months. As a result, customers received quotes that were 15% lower than the current list price. Over the course of several weeks, the company lost more than $250,000 in revenue before anyone noticed the discrepancy. What started as a simple data quality issue escalated into reduced revenue, higher costs to correct contracts, and strained customer relationships.

Sensitivity risk

Agents may process and share sensitive data without appropriate controls based on context and intended use. The critical issue isn't just identifying sensitive data types (PII/PHI/PCI/PSI), it's determining whether data can be shared and how. Data classification for sharing determines if information should be Public, Private, Internal, Restricted, or Confidential. Without proper classification systems that account for data sensitivity levels and appropriate usage contexts, organizations face regulatory fines and compliance violations when AI systems inappropriately share or expose data that should have restricted access. AI can't distinguish by itself, between data that's safe to share internally versus data that requires restricted access or redaction.

Example: An AI chatbot handling HR requests began including confidential salary information in responses to general org chart questions. This led to compensation details being shared more widely than intended, prompting the company to pause the tool and review its access controls. The issue resulted in added compliance work, increased legal review costs, and the need to rebuild employee trust in how sensitive information is managed.

Certification risk

Data is used "as-is", without validation that it's appropriate, accurate, or authorized for AI use. Agents may pull from test datasets, deprecated sources, or data that was never validated or certified for automated decision-making. Teams may not discover these issues until mistakes make their way into production or customer-facing outputs.

Example: A marketing AI pulled data from a test environment that had never been certified for production use. As a result, fake test accounts were included in customer segments, leading to wasted email sends, inflated campaign metrics, and added operational cost to clean up the lists.

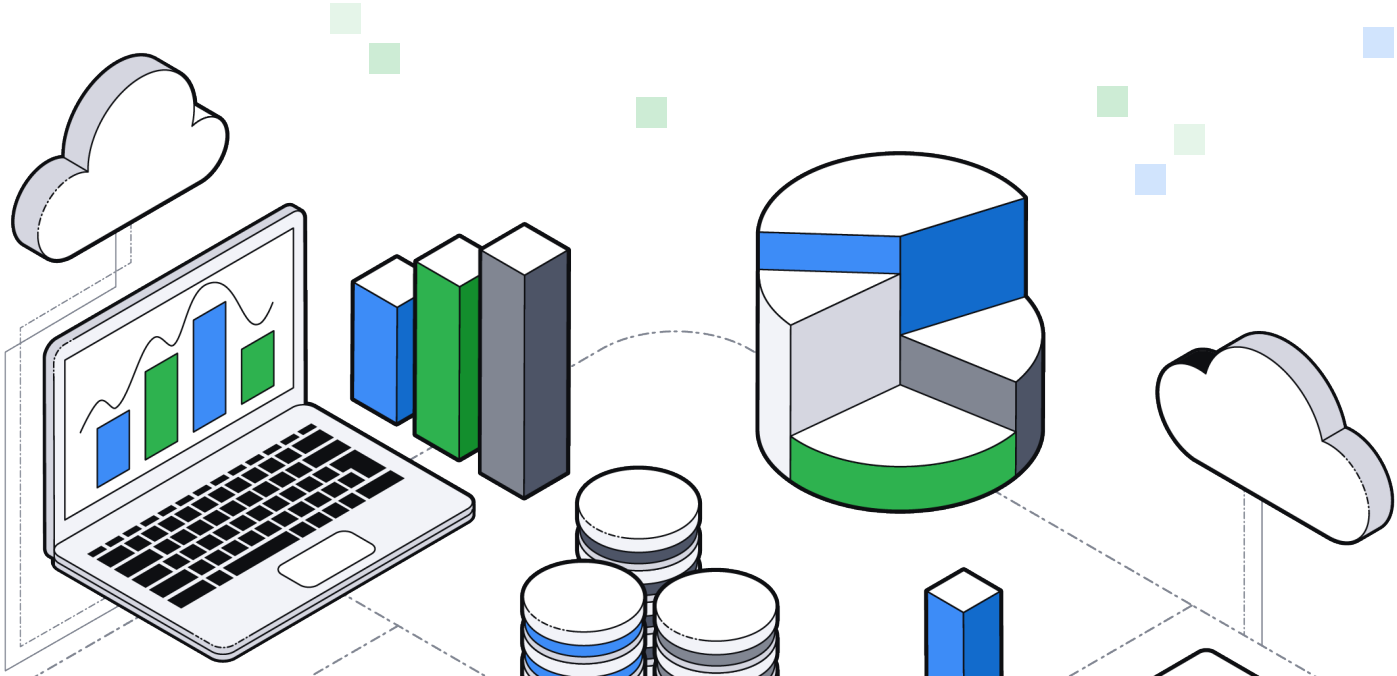

How the most mature orgs are operating

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:

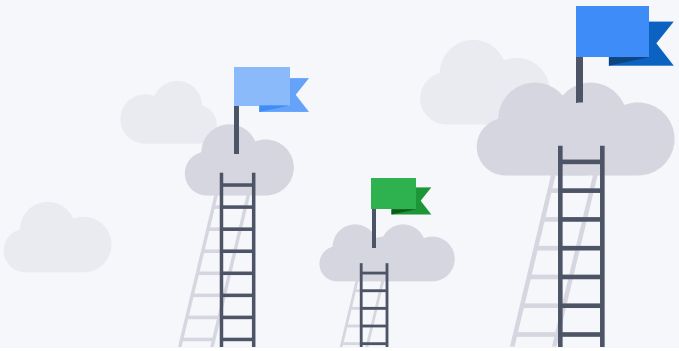

Quality Foundation: Know Your Data

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 Unaware to Aware

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

Quality Foundation

» Start Here

Create a list of data objects used by AI
Inventory which datasets, tables, and sources your AI systems are accessing

Define what columns can be used by AI
Establish column-level specifications for what data fields should be available

---

Sensitivity Controls

» Build Visibility

Create and make available a list of data objects to be used by AI
Publish an approved inventory of datasets authorized for AI access

---

Certification Framework

» Establish Oversight

Document if data is approved to be used by AI
Create formal records of which data has been validated for AI use

Document if process management and oversight exist for AI data
Establish basic management processes aligned with standards like ISO-9001 and CMMI

Why these five? These steps create the foundation of visibility and basic control that everything else builds on. Once you have clear inventories, defined usage parameters, and documented approval processes, you can systematically work through the remaining capabilities to reach full operational maturity.

Where to start

Here's how teams at your stage are beginning to close the gap:

Inventory what you know

Start with your agent footprint. Which systems or prototypes are querying internal data today? What can they access? Who owns that data? But don't stop at just cataloging what exists — you need to understand what should and shouldn't be available. Ask: Is this data appropriate for AI use? Does it meet quality standards for automated decision-making? Has it been validated or certified for agent access? If you can't answer these questions, visibility is your first blocker.

Define trust signals

What *should* qualify a dataset for AI use? Think about freshness, ownership, certification status, and sensitivity level. But also consider business fit — is this data appropriate for the decisions your AI will make? Even a basic rubric that covers quality thresholds, certification requirements, and appropriate use cases will help create alignment between teams.

Don't rely on intention

If your controls are based on "we trust our teams to do the right thing," you're missing the oversight layer. AI agents don't interpret intent — they need guardrails. This includes preventing access to data that shouldn't be used in AI, not just managing how approved data gets used.

Socialize the risk

Share this assessment with security, data, and platform stakeholders. Managing data for AI agent usage isn't a side task. It's foundational infrastructure. And everyone's going to feel the pain if it's missing.

Why it matters now

The speed of AI development is outpacing most organizations' ability to govern it. Teams that stay at this stage typically face one of two outcomes:

Public failure - Sensitive data gets exposed, financial impact occurs (like fines or penalties), or you're forced into an embarrassing walkback

Quiet stall - Security, legal, or privacy teams block or defund AI efforts because they don't trust what's happening under the hood

You don't need a perfect system today, but you do need visibility into what your AI is actually doing.

The organizations that build this foundation 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 already exists.

Most AI initiatives don't fail because the models are wrong. They fail because the data wasn't ready, and no one knew until it was too late.

What to do next

Ready to build your AI Trust foundation but need guidance?

Moving from Unaware to Aware requires establishing new processes, inventorying existing systems, and creating documentation that didn't exist before. 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