Bigeye

AI Trust Assessment Results

Managed Stage

You have comprehensive processes, but they can't match AI's speed.

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 ARE HERE)

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 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 Managed stage of AI trust maturity.

At this stage, your organization has implemented structured data classification, quality monitoring, and approval processes across much of your data landscape. You have comprehensive governance frameworks in place - including AI governance councils, cross-functional forums, and systematic risk management processes. Your teams understand which data is validated for AI use, and you have consistent controls guiding agent behavior.

You monitor data quality across multiple dimensions with automation, track privacy risks systematically, and maintain detailed records of training datasets and versions. Most high-impact datasets go through defined validation processes, and you have searchable catalogs that business teams actively use. Issues are assigned to documented owners through established processes.

However, while your processes are comprehensive, they're not yet fully automated or integrated across all systems. You may have some gaps in coverage, and your controls may still require manual intervention rather than operating with full real-time enforcement.

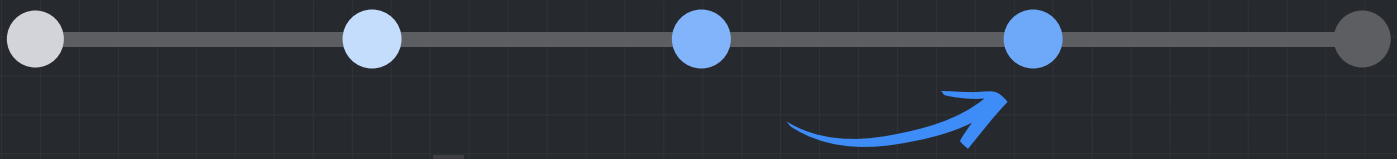

The real cost of staying here

While you have sophisticated processes in place, staying at the Managed level creates specific limitations as your AI ambitions grow:

Quality risk

You regularly measure quality dimensions with automation, but may still have gaps in comprehensive, real-time monitoring. Your validation processes work well for high-impact datasets, but as AI usage expands to more diverse use cases, manual validation steps can become bottlenecks. Without fully automated quality assurance integrated into all workflows, quality issues might still slip through in edge cases or less-monitored datasets.

Example: A logistics AI performed well on validated existing supply chain data, but when the company expanded into new regions, the monitoring didn't extend automatically to those local datasets. The AI made routing decisions based on outdated road infrastructure, resulting in delivery delays for 14% of regional orders. Customer satisfaction scores in those areas dropped by 10 points, creating reputational risk even though the core system was sound.

Sensitivity risk

You have structured privacy risk management and systematic data classification, but without fully automated, policy-driven enforcement, sensitive data exposure can still occur during rapid AI deployment cycles. Your scanning tools work well for most datasets, but may not catch all sensitive data in real-time as new data sources are added or existing data changes.

Example: A customer analytics AI accessed newly integrated acquisition data that contained PII from a different privacy jurisdiction. The systematic review processes caught this eventually, but not before the AI had processed several batches of data that should have been restricted under the acquired company's data governance policies.While no breach occurred, the company had to conduct a formal compliance audit and issue disclosures to regulators, consuming three weeks of legal and compliance resources.

Certification risk

You have comprehensive approval processes and detailed dataset tracking, but without complete automation, certification can become a scaling bottleneck. Your governance frameworks work well for current AI deployments, but may not be able to keep pace with rapid scaling or real-time decision-making requirements.

Example: A financial services AI needed to incorporate new market data sources during volatile trading periods. The systematic approval process worked correctly but took several days to validate the new data sources, causing the AI to make decisions on slightly outdated information during critical market movements. The firm underperformed compared to competitors who acted on fresher data.

The fundamental challenge at the Managed stage is that manual oversight, even when systematic and comprehensive, cannot match the speed and scale requirements of fully operational AI systems.

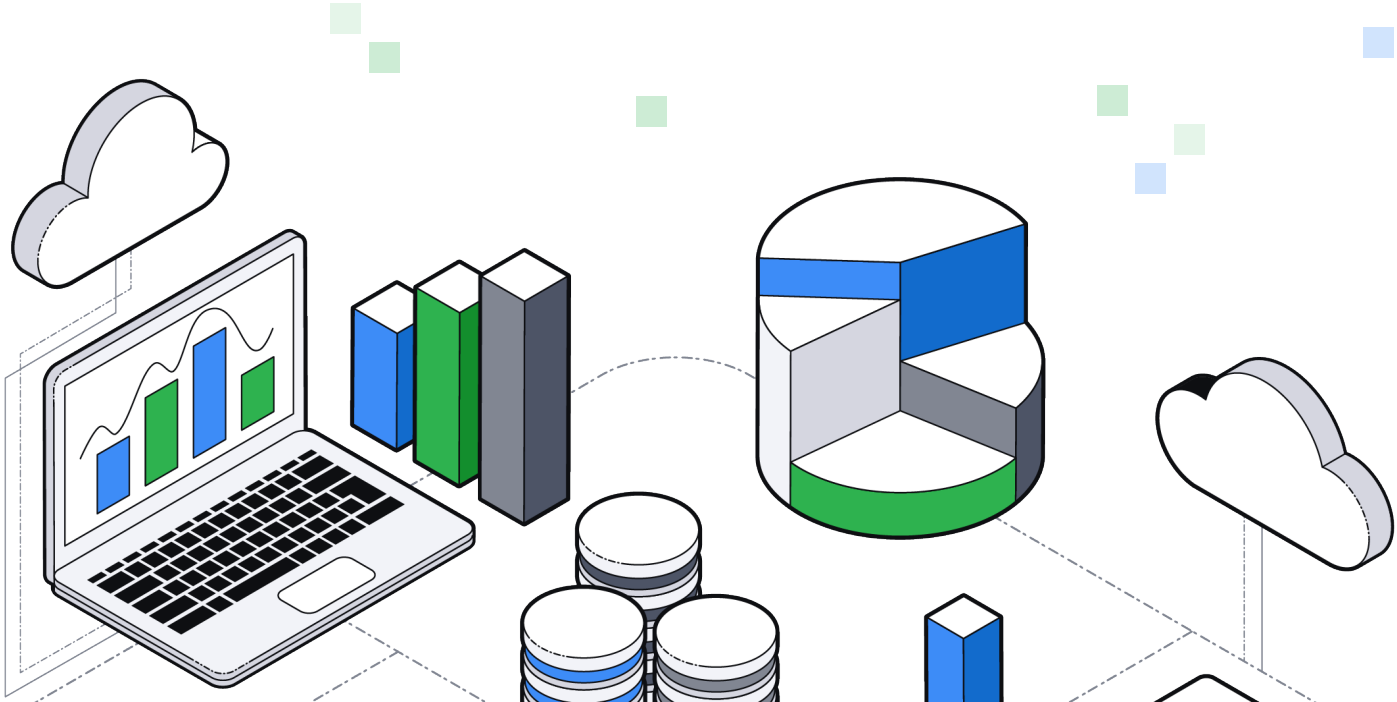

How the most mature orgs are operating

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:

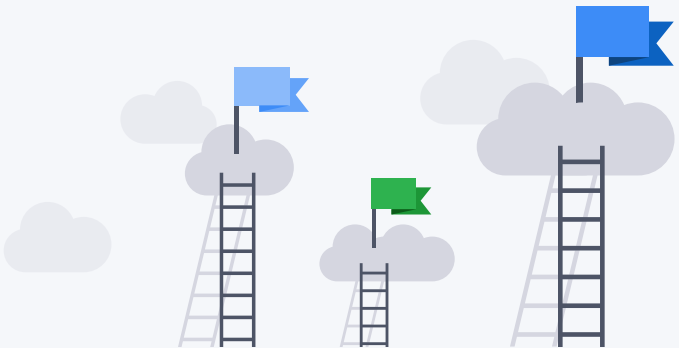

Quality Foundation: Know Your Data

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 Manged to Operational

While the above represents the full scope of AI Trust maturity, you don't need to rebuild everything at once. To move from Managed to Operational, focus on automating and integrating your existing comprehensive processes:

Quality Foundation

Achieve Full Automation

Implement automated quality assurance

Build on your existing quality monitoring by integrating automated validation into all AI workflows with real-time alerts and blocking capabilities

Expand comprehensive coverage

Ensure your quality processes cover all datasets and use cases, not just high-impact ones

Sensitivity Controls

Automate Policy Enforcement

Implement policy-driven classification

Automate your existing classification processes with real-time enforcement tied to data use and access

Achieve comprehensive scanning

Extend your current scanning tools to provide real-time detection across all AI data workflows

Certification Framework

Integrate All Systems

Automate certification tracking

Build on your detailed record-keeping by implementing automated lineage tracking that captures all datasets, versions, and changes throughout model lifecycles

Integrate comprehensive oversight

Connect your various governance frameworks into seamless, automated oversight with real-time monitoring and response

What you already have: As a Managed organization, you have sophisticated, systematic processes across all three areas. Your challenge now is automation and integration - making these processes operate seamlessly at machine speed without manual intervention.

Where to start

Here's how teams at your stage achieve full operational maturity:

Automate your quality validation

Build on your existing quality monitoring by implementing automated validation that can operate in real-time across all AI workflows. Convert your current approval processes into automated gates that can validate data fitness without manual intervention.

Implement real-time policy enforcement

Extend your systematic data classification by implementing automated policy enforcement. Your current structured processes should be converted into automated rules that can enforce access controls and usage policies in real-time.

Integrate your governance frameworks

Connect your various oversight mechanisms - AI governance councils, risk management processes, and approval workflows - into integrated systems that can operate automatically while maintaining the comprehensive oversight you've built.

Expand automation coverage

Ensure your automated processes cover all AI use cases, not just high-impact ones. Scale your current systematic processes to operate automatically across your entire AI data landscape.

Implement end-to-end observability

Build on your current monitoring and tracking by implementing comprehensive observability that provides real-time visibility into all aspects of AI data usage, from pipeline performance to policy compliance.

Why it matters now

Organizations that stay in the Managed stage often face a specific challenge: their comprehensive processes work well but cannot match the speed and scale demands of advanced AI deployments. This creates two key limitations:

Manual processes become scaling bottlenecks - Your systematic validation and approval processes work excellently for current AI deployments, but manual elements become constraints as AI usage accelerates or requires real-time decision-making capabilities.

Governance lag creates risk windows - Your comprehensive oversight frameworks provide excellent governance, but the time required for manual review and approval can create windows where AI systems operate without complete validation, especially during rapid deployment cycles.

The organizations that successfully transition to Operational stage see dramatic improvements in their ability to scale AI confidently. They can deploy new AI use cases rapidly because validation happens automatically, respond to changing requirements in real-time because policies are enforced automatically, and maintain comprehensive governance because oversight operates seamlessly at machine speed.

Most importantly, reaching the Operational stage means your AI initiatives can scale without trading off governance for speed - your comprehensive processes operate automatically, providing both the rigor you've built and the agility advanced AI requires.

What to do next

Ready to complete your systematic frameworks?

Moving from Managed to Operational requires transforming your systematic frameworks into fully automated, integrated systems that can operate at machine speed. Many teams find this transition 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