ACTIAN™
a division of HCLSoftware

The Bad AI Prevention Checklist

Discover, Trust, and Activate Your Data

Follow six steps to ensure you have the right data, context, and guardrails to keep AI accurate, explainable, compliant, and reliably aligned to business outcomes.

Before AI can reliably support your use cases, such as enabling AI agents that autonomously analyze data and take actions, you need confidence in the data behind every answer and action. Most "bad AI" outcomes come back to the same root cause: disconnected, low-quality, or poorly governed data.

The gap between success and failure isn't about your model choice. It's about whether you can discover the right data, trust its quality, and activate it safely at scale. According to Gartner®[1]:

Organizations

63%

are unsure or don't have the right data management practices for AI.

AI Projects

60%

will be abandoned by 2026 because they aren't supported by AI-ready data.

That means most organizations will struggle not because their AI is weak, but because their data foundations are. This checklist helps close that gap.

1 Gartner, **"Lack of AI-Ready Data Puts AI Projects at Risk,"** February 2025.

The Path to Confident AI

Discover

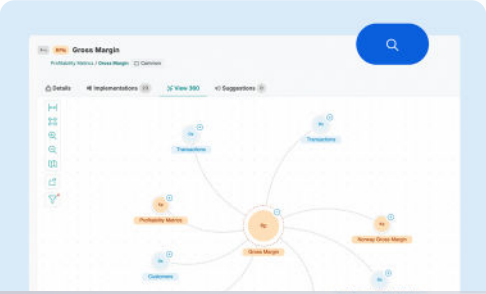

- Find, understand, and access data instantly.
- Eliminate 80% of time spent searching for and prepping data.

Example: Surfacing all customer interaction data for a churn model.

Trust

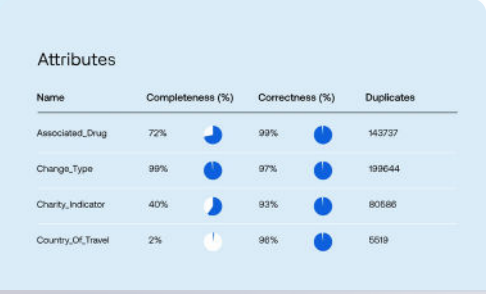

- Build trust, mitigate risk, and ensure compliance by design.
- Pass audits without slowing AI innovation.

Example: Knowing PII is masked before training an LLM.

Activate

- Activate data with agents and automation.
- Ensure AI success with production-ready data.

Example: Publishing a certified data product to your AI agent framework.

ACTIAN™
a division of HCLSoftware

6 Steps to Discover, Trust, and Activate Data for AI

Ensure data quality, readiness, and confidence at every point along your AI journey. These practical steps will help you deploy AI agents and other use cases with faster speed and less risk.

#1: Align stakeholders around a shared AI data strategy

To avoid bad AI surprises, every team must understand how they contribute to discovering, trusting, and activating data. Use these questions to define expectations for each role:

Business Users

- [ ] What AI use cases matter most to our customers and business?
- [ ] Which data do we actually need and not need to support them?
- [ ] Where could AI outputs create risk if the data is incomplete, biased, or outdated?

Action: Define specific outcomes and identify the minimal, most relevant datasets to support them. Avoid the "collect everything" mentality that adds complexity without adding value.

Data Stewards

- [ ] How will we define and enforce policies for data access, governance, and use for AI?
- [ ] How will we protect sensitive data while still enabling powerful AI use cases?
- [ ] What is our approval process for new data sources before they are used by models or AI agents?

Action: Work with IT, data engineers, and business teams to define guardrails for how data can be used in different AI scenarios. Document policies and enforce them through automated controls, not just manual reviews.

IT Developers / Data Engineers

- [ ] How will we integrate and transform data so it's discoverable, governed, and ready for activation?
- [ ] Where will we centralize metadata, lineage, and business context?
- [ ] How will we operationalize data products and contracts for reuse across AI initiatives?

Action: Design pipelines and data products that share common rules for validation, governance, and performance. Automate as much as possible to reduce manual work and accelerate delivery.

Business Analysts / Data Scientists

- [ ] How can we ensure self-service access to trusted, well-defined data?
- [ ] How do we know that KPIs, metrics, and definitions are consistent across dashboards, models, and agents?
- [ ] How will we choose and maintain training data and document its limitations?

Action: Standardize definitions, build shared business glossaries, and collaborate with domain experts to ensure models and metrics are grounded in agreed-upon data.

Business benefit: When business, data, and IT stakeholders align on how they discover, trust, and activate data, you reduce blind spots, accelerate time to value, and lower the risk of bad AI outcomes.

#2: Set shared standards for trustworthy data

AI amplifies whatever data you feed it, whether that information is high or low quality. That means you need clear quality and governance standards so teams can trust the data that powers models, agents, and applications. Ask critical questions to strengthen your foundation for trusted data:

- [ ] How much data, and what type of data, is needed for priority use cases?

Action: Size and scope data assets now to avoid scale or performance issues later. Avoid over-collecting data that increases risk or cost without improving model performance.

- [ ] How will we classify data, such as public, internal, restricted, and regulated, based on sensitivity, the criticality of how essential it is to operations and outcomes, and allowed uses?

Action: Tag data so you know what can be safely used for experimentation, what requires strict controls, and what cannot be used for AI.

- [ ] How will we set and enforce accuracy and completeness thresholds?

Action: Define clear tolerances for data quality metrics. Use them as Go/No-go criteria before feeding data into models or AI agents.

- [ ] How will we validate and monitor external or third-party data?

Action: Treat external datasets as "guilty until proven trustworthy." Establish checks to assess reliability and bias, and document how those sources are used.

- [ ] How will we prevent poor-quality data from entering our systems?

Action: Push validation and data quality controls as far upstream as possible so you're not constantly fixing issues downstream in AI workflows.

Business benefit: Clear standards for trustworthy data reduce rework, limit hallucinations, and create a measurable foundation for responsible AI.

#3: Manage data quality as a lifecycle, not a one-time project

AI use cases, regulations, and data sources evolve constantly. Trust and quality must be managed as a lifecycle with ongoing visibility, not a one-time deployment. For example, when product return policies update, downstream AI agents need fresh training data to avoid inaccurate refund decisions. These questions help establish data quality standards:

- [ ] How will our use cases and required data evolve over time?

Action: Build a roadmap that anticipates new data sources, changing regulations, and evolving use cases. Design data pipelines to flex with current and emerging requirements.

- [ ] How will we extend data quality and governance frameworks to existing pipelines?

Action: Add data observability checkpoints, policy enforcement, and lineage into current data flows, not just new ones.

- [ ] How will we test and validate data flows before they power AI experiences?

Action: Use controlled environments for staging and simulation. Combine automated tests with human review before promoting changes to production.

- [ ] How will we continuously monitor and measure data quality?

Action: Implement a real-time dashboard with alerts tied to key data metrics such as accuracy, timeliness, and completeness so issues can be detected and resolved early, before they impact outputs. Add observability to at least one existing pipeline before expanding widely because it helps demonstrate ROI quickly.

Business benefit: Managing data trust as a lifecycle, with observability, governance, and feedback loops baked in, creates a durable foundation for AI that can adapt as your business and regulatory environments change.

ACTIAN™
a division of HCLSoftware

#4: Create an activation strategy that balances risk and speed

Every organization must decide how fast to move with AI and how much risk they're willing to take. A clear activation strategy keeps innovation and guardrails in balance. Here are some considerations:

- [ ] Have we defined a safe "innovation zone" for experimentation versus production? This can be an employee-facing prototype, internal tooling, or R&D sandbox.

Action: Separate experimental use cases from mission-critical ones. Apply lighter controls in sandbox environments and stricter governance where AI touches customers, finances, or regulated data.

- [ ] How will we determine which data and use cases are ready to move from experiment to production?

Action: Set explicit criteria for promoting pilots, such as performance metrics, data quality scores, and governance approvals.

- [ ] Can we prevent bad data from driving high-risk AI decisions?

Action: Establish quality checks and automated validations before data is activated in critical workflows.

- [ ] How will we define, publish, and manage data contracts for AI use cases? A data contract defines the schema, quality thresholds, update frequency, and ownership model for a data asset.

Action: Define who owns the data, how it can be used, and what guarantees apply to quality and freshness. Treat data products and contracts as first-class assets.

- [ ] Are we prepared for changing regulations and AI-specific guidance?

Action: Choose an architecture and platform that can adapt quickly to new rules. Favor approaches where policy updates can be enforced centrally and managed automatically. Consider tiering AI use cases as low, medium, or high risk to apply different governance rules.

Business benefit: A deliberate activation strategy helps you move faster with AI where it's safe to do so while minimizing the risk of costly errors, compliance violations, and reputational damage.

#5: Use data intelligence to discover and activate the right data for AI

Data volume alone doesn't create value. You need intelligence, such as metadata, lineage, and business context, to find the right data, understand it, and safely activate it for AI. Asking these questions will help:

- [ ] Can teams quickly discover the data they need?

Action: Use a unified data intelligence platform to make data assets easily searchable, accessible, and explorable.

- [ ] Do we have lineage that shows where data originates and how it's transformed?

Action: Implement data lineage that's understandable to both technical and business users. Make it easy to trace changes back to their sources.

- [ ] Are we enriching data with business context, not just technical metadata?

Action: Maintain glossaries, classifications, and relationships that align with how the business actually talks about customers, products, locations, and risks.

- [ ] Do we make data intelligence accessible to the teams that need it?

Action: Provide self-service data access with appropriate guardrails so developers, analysts, and business users can confidently use trusted assets.

Business benefit: Data intelligence turns fragmented data into discoverable, understandable, and reusable assets that power AI responsibly. Teams spend less time searching for data and more time delivering value.

#6: Make trust measurable with data observability

To prevent "when bad AI happens" moments, you need real-time visibility into how data behaves as it flows into and through AI systems. For example, spotting and mitigating data drift can prevent a model from suddenly misclassifying high-value customers as low-value. Trust your data by asking:

- [ ] Can we monitor data quality and pipeline health in real time?

Action: Deploy observability, such as signals and metrics, that continuously evaluates quality, drift, and anomalies, rather than relying on occasional batch checks or sampling.

- [ ] How quickly can we detect and resolve data issues before they affect AI outputs?

Action: Create policies and automate workflows for alerting, triage, and remediation. Tie them to service-level objectives so teams know how quickly to respond.

- [ ] Do our observability capabilities cover hybrid and multi-cloud environments?

Action: Choose tools and architectures that provide unified visibility across on-premises, cloud, and hybrid environments to avoid blind spots.

- [ ] Are we relying on manual monitoring and reactive fixes?

Action: Shift to AI-driven monitoring that predicts disruptions before they occur, reduces repetitive work, and lets teams focus on higher-value innovation. Start by automating anomaly detection on one high-impact table or pipeline.

Business benefit: Modern data observability makes trust visible and actionable. It reduces the risk of undetected issues fueling bad AI outcomes and helps you protect both SLAs and stakeholder confidence.

Don't Just Build AI. Build It on Trusted Data Intelligence.

AI is only as reliable as the data and guardrails behind it. Success isn't about collecting more data or chasing the latest model. It's about being able to discover the right data, trust its quality and governance, and activate it with confidence.

By using this checklist, your organization can move beyond experimentation to take measurable steps toward AI you can explain and trust. With the right strategy, collaboration, and data intelligence foundation, you'll be positioned for sustainable innovation.

About Actian

Actian empowers enterprises to confidently manage and govern data at scale. Organizations trust Actian data management and data intelligence solutions to streamline complex data environments and accelerate the delivery of AI-ready data. Designed to be flexible, Actian solutions integrate seamlessly and perform reliably across on-premises, cloud and hybrid environments. Learn more about Actian, the data and AI division of HCLSoftware, at actian.com.

Tel +1 512 231 6000 Fax +1.512.231.6010 710 Hesters Crossing Road, Suite 250, Round Rock, TX 78681 **Actian.com**

ACTIAN™
a division of HCLSoftware