cloudfactory

AI TRUST

CLAIM 4471 · CONFIDENCE 0.62 · ROUTED TO EVALUATOR



SOLUTION BRIEF

Trusted AI at scale.

Keep production AI accurate, accountable, and audit-ready, from the first dataset to the millionth decision.

CRITICAL SLA TIER

1 to 5 min

Response to low-confidence and high-risk outputs, inside the live workflow

TRUSTED PARTNER

700+

of the world's most ambitious AI companies, since 2010

01 · THE TRUST GAP

WHERE TRUST BREAKS

A pattern many enterprises recognize: a property and casualty insurer runs millions of dollars in claims decisions through AI every day.

- **Data failure**
  Ungoverned, incomplete data enters the pipeline.
- **Model output**
  A probabilistic guess, presented as 87% certainty.
- **Human override**
  A person overrules the model and documents nothing.
- **Audit failure**
  No system of record, nothing to reproduce.
- **Enterprise risk**
  Regulatory exposure, financial liability, and AI the business can't deploy.

THE COST OF A SILENT ERROR

Months later, no one can explain how the call was made. Then the regulator asks.

The model was never the weak point. The missing record was.

BUILT FOR HIGH-CONSEQUENCE WORK

- Insurance claims and underwriting
- Transportation and logistics
- Robotics and manufacturing
- Energy, oil, and gas
- Finance operations

WHY PRODUCTION AI STALLS

AI breaks at the seams, not in the model.

Enterprises have proven AI works in a pilot. Scaling it into decisions that carry financial, safety, and regulatory weight is a different problem: one of confidence, accountability, and proof.

The trust gap isn't a single failure. It's a systemic gap across data, models, oversight, and integration, and it compounds as volume grows. A model that's right 95% of the time still makes thousands of wrong calls at enterprise scale, and most of them look exactly like the right ones.

Closing that gap takes more than better models. It takes a system that knows which outputs to trust, routes the rest to qualified people, records every correction, and feeds what it learns back into the model so the same error doesn't happen twice.

Four gaps that keep AI out of production

01 / DATA

Messy, unstructured data

Siloed, inconsistent, and rarely AI-ready. Without a clean semantic foundation, models learn the wrong things.

02 / MODELS

Generic foundation models

Trained on public data, not your context, so they lack the domain knowledge reliable outcomes depend on.

03 / OVERSIGHT

No human in the loop

With no expert check between model output and business decision, automated pipelines let errors compound.

04 / INTEGRATION

Fragmented integration

Models that perform in isolation still fail when outputs flow into real enterprise workflows.

A VISUAL LANGUAGE OF CONFIDENCE

Every output carries a score.

Confidence is shown, not hidden. Each state has one route and one owner.

0.97 · passed	Goes straight through to the model and the business
0.62 · in review	Routed to a specialist agent or an Evaluator
0.31 · held	Contained and held back from the model

THE GOAL

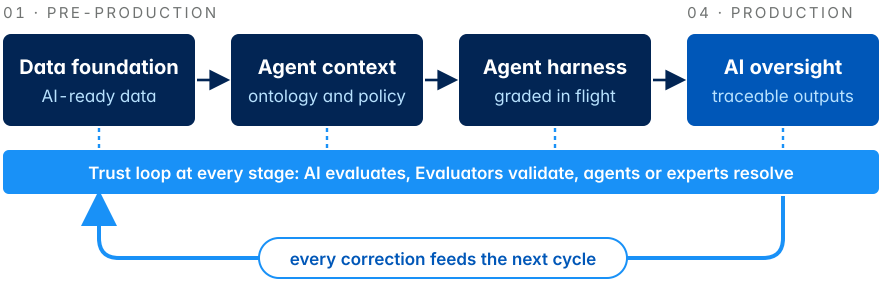
The aim isn't AI that's never wrong. It's AI where being wrong is visible, contained, and corrected before it reaches the business.



HOW CLOUDFACTORY CLOSES THE GAP

One loop, from data to trust.

Enterprise data becomes AI-ready. Context is validated by experts. Every agent step is graded in flight, so low-confidence decisions reach an Evaluator before they act. Oversight then keeps every production output traceable.



The trust loop starts by knowing when AI is wrong.

Most organizations can't tell when their AI is wrong, so they can't manage the risk of deploying it. The trust loop finds, validates, and resolves errors at scale.

01
Automated error detection
AI evaluates your AI against your ontology, policies, and rules to flag likely errors and non-compliance.

02
Intelligent routing
CloudFactory's AI Platform sends each flagged issue to the Evaluator with the right expertise to confirm whether it's a genuine problem.

03
Validation and remediation
Once an error is confirmed, corrective action is orchestrated through machines or people, whichever fits.

04
Continuous feedback
Outcomes feed back into the system, so detection and handling improve with every cycle.

Steps 01 to 04 are logged and auditable: every detection, route, decision, and fix is on record.

Automated error and compliance detection, human validation only where it's needed, and orchestrated remediation: continuous AI assurance.

FOUR PILLARS, ONE LOOP

Data Foundation 01
Built on data you can trust
AI does the work, a second AI layer scores it, and data analysts validate only where confidence falls short. Every modality, one pipeline.

Model & Agent Evaluation 02
Proof it works, before it ships
Evaluation against your definitions and production data, with trace-level grading of agent reasoning, tool calls, and decision paths.

AI Trust 03
Catch what it gets wrong, prove you caught it
Policies become enforceable guardrails, with full lineage across prompt, inference, tool call, and correction.

Managed AI Operations 04
We run it, at your SLA
Hosting, pipelines, and retrieval run together, with coverage up to 24×7. Drift routes to a person, not an alert queue.

Human expertise is a confidence mechanism built into the system, not a labor pool bolted on.

Ground truth for a robot that sees from the ground.

Syngenta's Robodog is a four-legged robot that carries depth cameras and an onboard GPU, running computer vision models at the edge in real-world environments.

Static training data can't keep pace with a robot moving through the real world. Robodog needed intelligence built for its own height and angle, plus an evolving analysis of the edge cases it meets in the field, from touching objects to partial views, so the model learns to detect and count instead of guessing.

CloudFactory also runs model-in-the-loop QA on Robodog's output, correcting missing, partial, and misplaced detections. The model predicts, a specialist corrects it, and each correction makes the next prediction better. It's the trust loop in miniature, running on physical AI.

3D

depth-camera imagery analyzed from the robot's own vantage point

3

defined model error types corrected in the loop: missing, partial, and misplaced detections

5

annotation methods matched to the task, from polygons to model QA

Three ways in. One accountable partner.

WAYS TO ENGAGE

PATH 01

AI Advisory & Engineering

Forward Deployed Engineers (FDEs) embed with your team to design, build, and ship, transferring knowledge as they go.

PATH 02

AI Applications

Packaged AI that delivers a business outcome from day one, led by our flagship, Invoice Intelligence.

PATH 03

CloudFactory's AI Platform

The software behind the loop, run by your own team, with CloudFactory Evaluators on call for exceptions.

Forward deployed to build. Managed to run. One accountable partner across both.

ABOUT CLOUDFACTORY

CloudFactory is an end-to-end AI trust partner, ensuring high-value AI processes and decisions run on data you can trust. Since 2010, CloudFactory has combined AI technology and human expertise to help enterprises and AI technology companies develop, deploy, and operate reliable AI across industries where trust is critical to operations.

CloudFactory's AI Platform, our FDEs, and our global team of CloudFactory Evaluators cover the full AI lifecycle, from Data Foundation, Model & Agent Evaluation, and AI Trust to Managed AI Operations. More than 700 of the world's most ambitious AI companies trust CloudFactory to keep their AI accurate in production.

START THE CONVERSATION

See the trust loop running on your workflows.

cloudfactory.com

Trusted AI at Scale.

© 2026 cloudfactory. All rights reserved.