TRUST3 AI BRIEF · AI APPS AND AGENTS

No AI ROI **without AI security**

How Trust3 AI fits the Build, Launch and Manage lifecycle described in Microsoft's *ROI Playbook for AI Apps and Agents*, so every new agent ships with the same discovery, access control and evidence as the last.

Trust3 AI

Discover every agent. Observe every decision.
Secure every action. Govern every data source.

INTRODUCTION

The ROI story, and **what it quietly depends on**

Microsoft's playbook makes a clear case. AI has moved from assisting people to acting for them, and the teams that turn agents into a repeatable production pattern are the ones that see returns. Teams that stay stuck in pilots pay for rework, incidents and rising run costs instead.

327%
three-year ROI reported for Microsoft Foundry[1]

<6 mo
payback period in the same study[1]

73%
of enterprises have no complete AI agent inventory[2]

3–10x
how far enterprises typically undercount their agents[3]

The argument in brief

The playbook's answer for lean teams is standardization. Set up the hard parts once on Microsoft Foundry and GitHub, then reuse them for every agent that follows. ROI becomes a function of cost-to-serve and time-to-value per workflow, and it improves as each new agent reuses the pattern.

It also says something security leaders will recognize: as agents gain autonomy, oversight has to grow with them, and governance added late is one of the main reasons agents stall between pilot and production.

Where Trust3 AI comes in

A repeatable delivery pattern only pays off if trust is repeatable too. Every agent needs an identity, an owner, access scoped to its purpose, a trace of what it did and evidence an auditor will accept.

Most enterprises also run agents beyond one platform: Bedrock, Databricks, Snowflake, Salesforce, coding agents in Cursor and Claude Code. Trust3 AI is the control plane for enterprise AI agents and data. Its Unified Trust Layer applies one policy model across every framework and cloud, including the agents you build on Foundry.

> No AI ROI without AI security. Standardize delivery once and ROI compounds; standardize security once and it stays compounded as agents, teams and platforms multiply.

[1] The Total Economic Impact™ of Microsoft Foundry, Forrester study commissioned by Microsoft, February 2026, as cited in Microsoft's ROI Playbook for AI Apps and Agents. [2] Trust3 AI survey of 437 technology experts. [3] Trust3 AI field data.

© 2026 Trust3 AI · trust3.ai

THE CHALLENGE

Where ROI leaks **after the pilot**

The playbook names four blockers that show up once an agent leaves the lab. Each one has a trust dimension, and each one gets more expensive with every additional agent.

Blocker in the playbook	How it shows up for security and data teams	What Trust3 AI adds
The repeat tax of fragmented tooling	Every agent gets its own access rules, logging and review. Controls drift apart and nobody can say which agent holds which entitlement.	One policy model, defined once and enforced across agents and 50+ data sources.
Customizing reliably as inputs change	New tools, new MCP servers and injected prompts change what an agent does after it was approved.	Continuous observation of scope drift, with a Trust Score that moves when behavior does.
Scaling operations without adding people	Manual reviews and ticket queues for access become the bottleneck as the agent count grows.	Just-in-time, auto-expiring grants and plain-English policy through GIA, the Governance Intelligence Agent.
Governance that arrives too late	Controls are bolted on per agent, or only cover agents on one platform.	Coverage from build time to audit time, across any framework and any cloud.

Three approval failures behind the numbers

The shadow agent

Built in an afternoon in a coding tool and wired to an unvetted MCP server. No ticket, no identity, no owner. An approval gate can't refuse what was never submitted.

The blank check

Access is static and role-based. One token opens everything the role can see, around the clock, with no grant tied to a purpose.

The moving target

Approval happens once, but agents change daily. The agent that was approved may not be the agent that is running today.

Each of these erodes the unit economics the playbook asks you to track. A shadow agent adds cost no one attributed. A blank check widens the blast radius of a single prompt injection. A moving target turns an audit into weeks of reconstruction.

© 2026 Trust3 AI · trust3.ai

THE LIFECYCLE

Build, Launch, Manage, **with trust at every stage**

The playbook organizes AI delivery into three stages. Trust3 AI follows the same stages and adds the controls that travel with each agent from its first commit to its last audit.

	BUILD	LAUNCH	MANAGE
MICROSOFT PLAYBOOK	Agentic DevOps with GitHub and Copilot; reusable templates; security in the developer workflow.	A consistent architecture on Foundry for models, orchestration, grounding, runtime and data.	Evaluation, observability, cost control and built-in governance as usage grows.
TRUST3 AI	**Discover.** Every agent, tool and MCP server found as it is built, including shadow AI, with an owner and identity attached.	**Score and scope.** A Trust Score before deploy, and purpose-based access to the data the agent is allowed to reach.	**Observe and secure.** Full traces, runtime guardrails, a kill switch and one-click audit evidence.

Discovered at build time → scored before deploy → enforced at runtime → replayable at audit

Two sides of one control plane

Agent Security

Agent DOS discovers, observes and secures agents wherever they are built: Foundry and Copilot Studio, Bedrock, Databricks, LangChain, Salesforce Agentforce and local coding agents. It covers the surfaces agents touch, including tools, MCP servers, APIs and SaaS apps.

Data Security

Purpose-based access control (PBAC) and fine-grained access control (FGAC) enforced natively at the source on Snowflake, Databricks, BigQuery, Iceberg and 50+ more, with no proxy hop and no added latency.

Every request carries the same three pieces of metadata: **agent identity**, **declared purpose** and **live policy state**. That is what lets one policy hold at the agent boundary and again at the data layer, even when one agent calls another.

© 2026 Trust3 AI · trust3.ai

BUILD DISCOVER FROM THE FIRST COMMIT

Build once, reuse everywhere, **including the controls**

The playbook's productivity engine is agentic DevOps: agents writing code, opening pull requests and shipping through CI/CD alongside developers. That speed is the point. It is also how agents appear faster than any registry can record them.

What changes at build time

Agent risk starts when a developer spins up an agent in Cursor or Claude Code, or attaches an MCP server nobody has reviewed. If discovery waits for a deployment ticket, the inventory is already wrong.

Trust3 AI scans developer environments, SaaS tenants such as Copilot Studio and Salesforce Agentforce, and cloud accounts through CloudTrail, platform APIs and SDK activity. Unregistered agents are flagged the first time they appear.

What you get

- A live inventory with owner, identity, connected tools and MCP servers, data access and status for every agent.
- A Trust Score from 0 to 10 for every discovered agent, weighted across policy compliance, scope adherence, security posture and behavioral baseline.
- Policy templates that become part of the reusable pattern, so the tenth agent starts with the same guardrails as the first.

How it fits alongside GitHub

GitHub Advanced Security catches secrets, vulnerable dependencies and code issues in the pull request. Trust3 AI answers a different question about the same change: which agent is this, who owns it, what will it be able to reach and for what purpose. Both belong in the loop before anything merges.

100%
of discovered agents scored, including shadow AI

300+
agents discovered during one audit at a Fortune 500 financial firm

10x
faster from proof of concept to production[4]

[4] Independently validated by Intellyx in a 2026 analysis of Trust3 AI production deployments.

© 2026 Trust3 AI · trust3.ai

LAUNCH SET THE BAR ONCE, MEET IT EVERY TIME

The minimum bar before you ship, **enforced rather than checked**

The playbook lists six essentials an agent needs before it goes live. They are the right list. The hard part is keeping them true for every agent, on every platform, after launch day. Here is how Trust3 AI turns each one into a control.

Pre-launch essential	Trust3 AI capability
Trusted sources	Unified catalog with semantic enrichment, lineage and classification across the data estate, so you know which sources an agent can ground on and what sensitive data they hold.
Clear ownership	Every agent in the inventory has a named owner and identity. Agents without one are flagged as shadow AI rather than allowed through.
Access rules	PBAC evaluated per request: the same identity can get different access for project A than for project B. Row, column and tag-based controls enforced natively, with zero standing access.
Grounding you can trace	Traces from prompt to retrieval to tool call to response, so any answer can be tied back to the data it came from and the policy that allowed it.
Basic monitoring	Real-time detectors for PII leakage, scope drift and behavioral anomalies, alongside the quality and latency signals your platform already collects.
Cost attribution	Every action is tied to an agent, an owner and a declared purpose, which gives token spend and usage a workflow to belong to.

Identity that travels through the whole chain

The playbook recommends unique agent identities and least-privilege access through Microsoft Entra Agent ID. Trust3 AI builds on identity infrastructure like this rather than replacing it. It carries the originating user's identity and the agent's purpose through every hop, including MCP and agent-to-agent (A2A) calls, down to the data source. Policy is then enforced at the agent boundary and again where the data lives.

That matters most where the data is not in one place. Grounding in OneLake and SharePoint can sit next to tables in Snowflake, Databricks and BigQuery. One policy covers all of them.

© 2026 Trust3 AI · trust3.ai

MANAGE KEEP ROI INTACT AS USAGE GROWS

Prevention at runtime, **evidence at audit time**

The playbook's final stage is about letting autonomy grow without losing control of quality, security and cost. For a security leader, that means stopping a bad action before it executes and proving what happened afterward.

Secure every action

- Runtime guardrails inspect each agent action and block high-risk tool and API calls before they fire.
- A kill switch on every agent, for the moment an indirect prompt injection or misconfiguration turns into an incident.
- Just-in-time grants and auto-expiring scopes, so an agent holds access only for the task in front of it.
- Native MCP and A2A protocol security, treating every MCP server as untrusted until proven otherwise.

Observe every decision

- Full-fidelity traces of the decision path, tool chain and data accessed, replayable on demand.
- A Trust Score that updates continuously, surfacing agents that drift from their approved scope.
- One-click evidence packs for the EU AI Act, NIST AI RMF, HIPAA, GDPR, SOX and more.
- GIA, the Governance Intelligence Agent, to query posture and define policy in plain English.

Working alongside the Microsoft stack

Foundry Control Plane, Purview, Azure Monitor and Defender for Cloud give you strong governance and protection for the agents and data inside the Microsoft estate. Trust3 AI extends the same discipline across everything else your agents reach, and adds purpose-based enforcement at the data source. Existing AI gateways keep their place too: Trust3 AI sits behind them and enriches every log with agent identity, purpose and the policy verdict.

84%
reduction in audit preparation time[4]

50+
data sources governed natively, with zero proxy latency

Dec 2027
EU AI Act high-risk obligations for most categories

[4] Independently validated by Intellyx in a 2026 analysis of Trust3 AI production deployments.

© 2026 Trust3 AI · trust3.ai

GETTING STARTED

Quick wins that **protect the first workflow**

The playbook advises starting with one high-volume workflow, such as ticket triage, invoice processing or document intake, then measuring how much faster the second workflow ships on the same pattern. Add trust to that first workflow and the second one inherits it.

1. **Find out what is already running**

   Connect Trust3 AI to your cloud accounts, SaaS tenants and developer environments. Most teams find agents they did not know about in the first scan. Start the inventory before you start the next build.

2. **Score the workflow you plan to scale**

   Give each agent in the target workflow a Trust Score and an owner. Fix the lowest-scoring issues, usually over-broad access and unvetted MCP servers, before you call it production.

3. **Replace roles with purpose**

   Move the workflow's data access from standing, role-based grants to purpose-based, just-in-time access. Define it once in plain English and reuse it for the next agent.

4. **Turn on runtime enforcement and evidence**

   Enable guardrails and the kill switch, then generate your first evidence pack. Track audit preparation time next to cost-to-serve and time-to-value.

5. **Reuse the pattern**

   When the second workflow launches, compare how long its security review takes against the first. That number is the trust equivalent of the playbook's time-to-value metric.

Proven at enterprise scale

40% of Trust3 AI customers are in the Fortune 500. A Fortune 100 fintech that hit AWS IAM limits replaced over-privileged roles with attribute-based access. A top-5 US healthcare retailer governs thousands of BigQuery tables across hundreds of GCP projects with no ticketing queue. A Fortune 500 financial firm became audit-ready without pausing operations.

© 2026 Trust3 AI · trust3.ai

NEXT STEPS

No AI ROI without AI security. **Build it into the pattern.**

The playbook shows how a repeatable delivery pattern turns AI into durable ROI. Trust3 AI makes sure that pattern carries discovery, purpose-based access, runtime enforcement and audit evidence with it, for every agent, across any framework and any cloud.

Built by the creators of Apache Ranger, Trust3 AI brings a decade of production policy enforcement to the agentic era.

Book a 30-minute live demo. Bring your stack. trust3.ai

Trust3 AI

This brief is published by Trust3 AI and refers to Microsoft's publicly available *ROI Playbook for AI Apps and Agents* (2026). It is not produced, sponsored or endorsed by Microsoft. Microsoft, Azure, Microsoft Foundry, Copilot Studio, Entra, Purview, Defender and GitHub are trademarks of the Microsoft group of companies. Other product names are trademarks of their respective owners. Third-party statistics are attributed to their original sources.