Trustible

How AI Governance Drives AI Outcomes

A deep-dive on AI risks and AI benefits

CHAI Leadership Summit 2026
Dana Point, CA

About Trustible

Who We Are

Trustible is the leading AI governance platform enabling safe and compliant AI adoption. Our AI governance platform enables enterprises to identify, measure, and mitigate AI risk to accelerate AI adoption. The company is headquartered in the Washington D.C. area. For more information, visit trustible.ai.

What We do

Responsible AI Governance Software & Solutions

Whether you are building internal AI models, leveraging generative AI or using AI-enabled SaaS applications, Trustible™ enables your organization to manage and mitigate AI risk, build trust, and accelerate responsible AI development. Plus, align your efforts to regulatory frameworks like the EU AI Act, NIST AI RMF, or ISO 42001.

Customer Snapshot

Percent of Trustible customers meeting these criteria

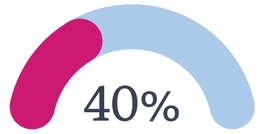

Fortune 500

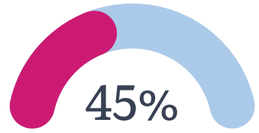

Healthcare

Operate Globally

Operationalizing CHAI's AI Governance Framework through Trustible

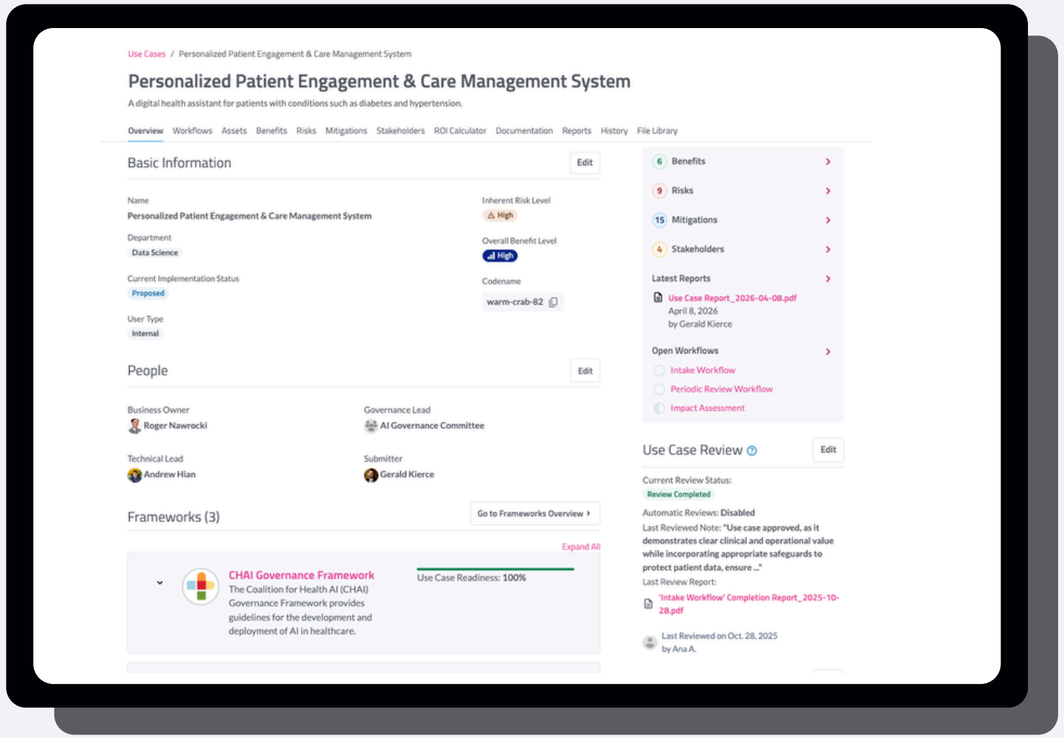

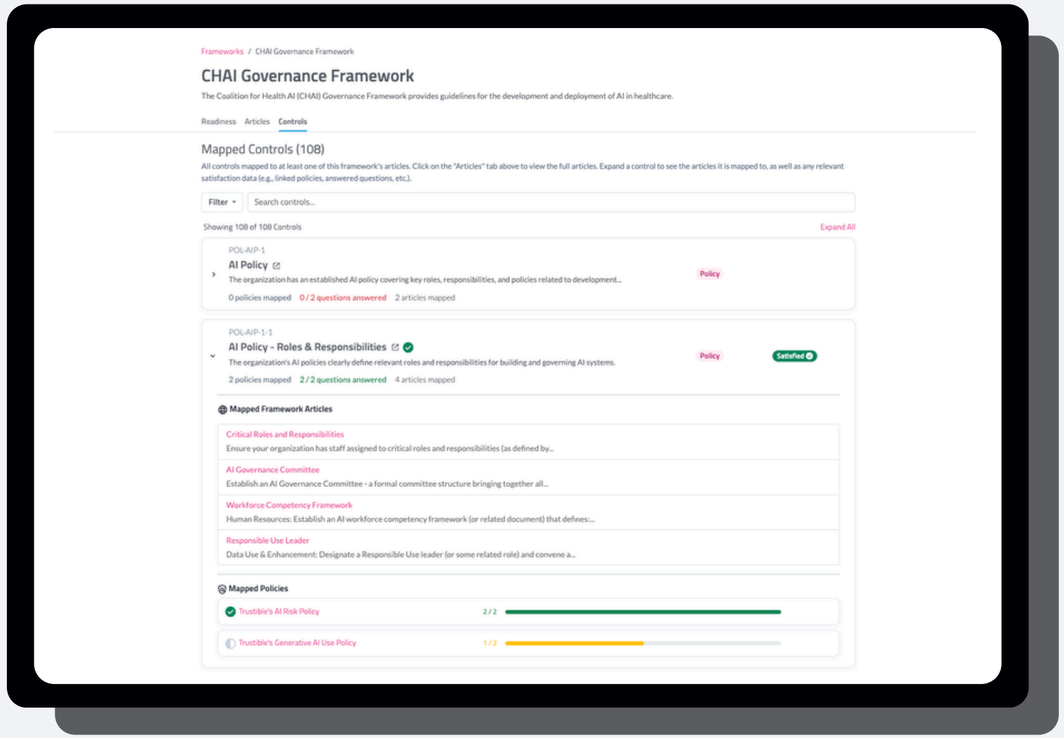

- Pre-built policies & controls framework
- Use case templates and model evaluations
- Healthcare AI risk & mitigation guidance

We likely don’t have to convince you since you’re here, but the metrics on investing in AI Governance are clear

3x

more likely to be an AI high performer

Organizations with strong governance leadership and oversight are 3x more likely to attribute >5% of EBIT to AI.

McKinsey State of AI, 2025

3.4x

more effective with governance platforms

Organizations deploying AI governance platforms achieve 3.4x higher governance effectiveness.

Gartner AI Governance, 2025

$4.4M

average loss from AI-related risks

99% of organizations reported financial losses from AI-related risks; 64% lost more than $1M.

EY Responsible AI Survey, 2025

The Problem

AI Is Moving Faster Than Governance Can Keep Up

Most organizations aren't limited by AI ambition.
They're limited by how fast governance can scale.

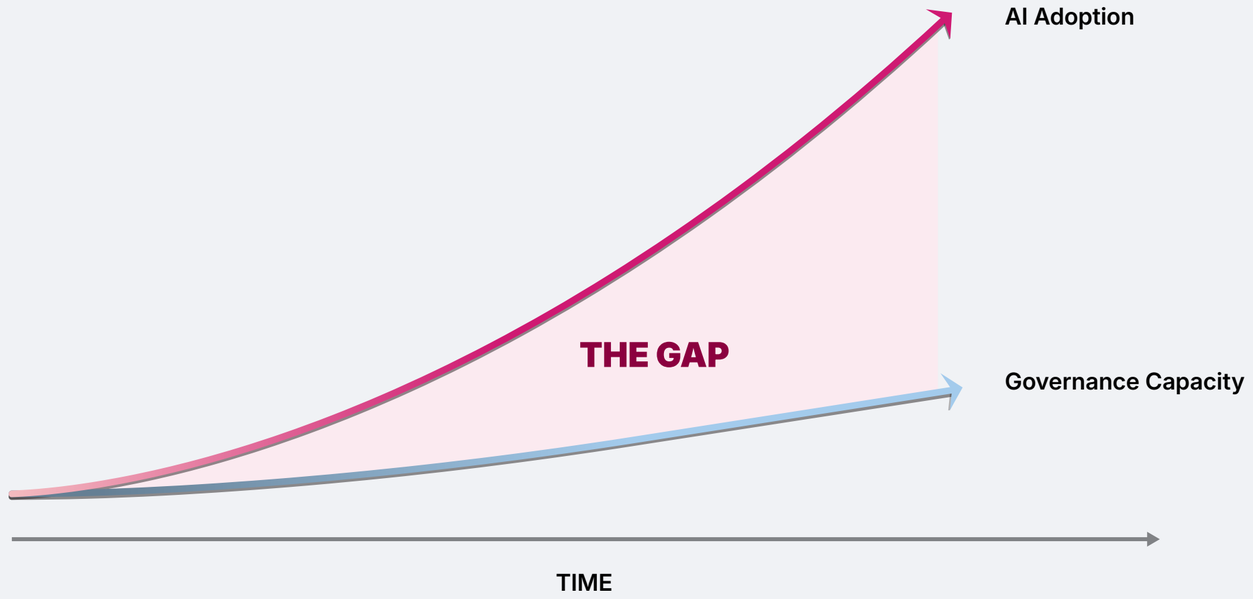

WHERE GOVERNANCE BREAKS DOWN

Intake Bottleneck

Reviews pile up.
Deployments stall.

Limited Visibility

Teams can't track what
AI exists across the org.

Regulations Outpace Controls

Rules change faster
than internal processes.

Lack of Trust

Customers and regulators
want evidence, not assurances.

This Conversation's Starting Point

What we assume you already have in place

- AI Policies are written
- AI Literacy training implemented
- Use case intake process established
- Initial risk / impact assessments done
- Governance committee roles defined
- Secure AI platforms created

This workshop will focus on two scenarios

Managing Risks & Harms of Agentic AI

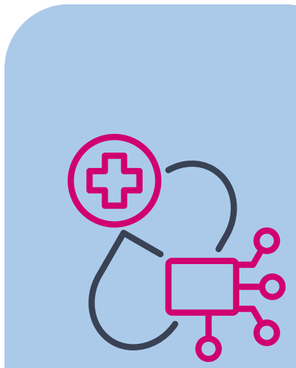

Measuring AI Benefits & Opportunities

Scenario 1

Managing Risks & Harms of Agentic AI

Measuring AI Benefits & Opportunities

Definitions

What is Agentic AI? What is an AI Agent?

The terms "agentic AI" and "AI agent" get used interchangeably, but they describe meaningfully different levels of autonomy — and from a governance standpoint, that difference matters. In practice, autonomy exists on a spectrum, but the two-level distinction gives teams a practical framework: agentic AI and AI agents present different risk profiles, require different oversight mechanisms, and raise different accountability questions.

Agentic AI

Human-triggered, AI-executed

- Human initiates the task
- AI determines how to execute
- Human typically reviews results

AI Agents

Independently triggered, autonomously executed

- Triggered by schedule or event
- AI sets sub-tasks and reasoning
- Minimal real-time human oversight

Mechanisms for autonomous decisions

How Agents take action

Tool Calling

Risk level: Bounded & Auditable

The AI calls pre-defined tools via APIs or MCP servers. The universe of possible actions is bounded by what tools exist. Administrators control which tools each agent can access, at what scope, and can revoke access quickly. Most manageable from a governance standpoint.

Computer Use

Risk level: Broad Action Space

The AI navigates screens, clicks, and types — like a human user. It can interact with any software visible on screen, whether or not an API exists. From the target system's perspective, agent actions may be indistinguishable from a legitimate human user.

Code Generation & Execution

Risk level: Open-Ended, Hardest to Audit

The AI writes novel code and executes it. That code could make network requests, manipulate files, query databases. The blast radius is bounded only by the runtime environment. Actions may blend into background system operations and be difficult to distinguish from legitimate automated processes.

Agentic AI presents different kinds of challenges

Autonomy
Agents take sequences of actions without human review at each step

Scope Creep
Agent goals can evolve in ways designers didn't anticipate

Third-Party Tool Use
Agents call external APIs, databases, and systems

Discussion

An AI agent is already making decisions in your organization. You found out today.

Scenario: A clinical department deployed an agentic AI tool that schedules follow-up appointments, routes care escalations, and sends patient messages without human review of each action. It has been running for six weeks. No governance intake was filed.

Considerations

- *How does your current intake process handle agentic AI that action rather than produces output for human review?*
- *Who is accountable for harm caused by an action the agent took — the department, IT, legal, or the vendor?*
- *What types of questions should your risk assessment ask to understand the agent's risk, access, triggers, etc?*
- *What is your threshold for requiring human approval before an agent acts, and is that written down?*

Scenario 2

Managing
Risks & Harms
of Agentic AI

**Measuring AI
Benefits &
Opportunities**

A useful framework for prioritizing AI use cases

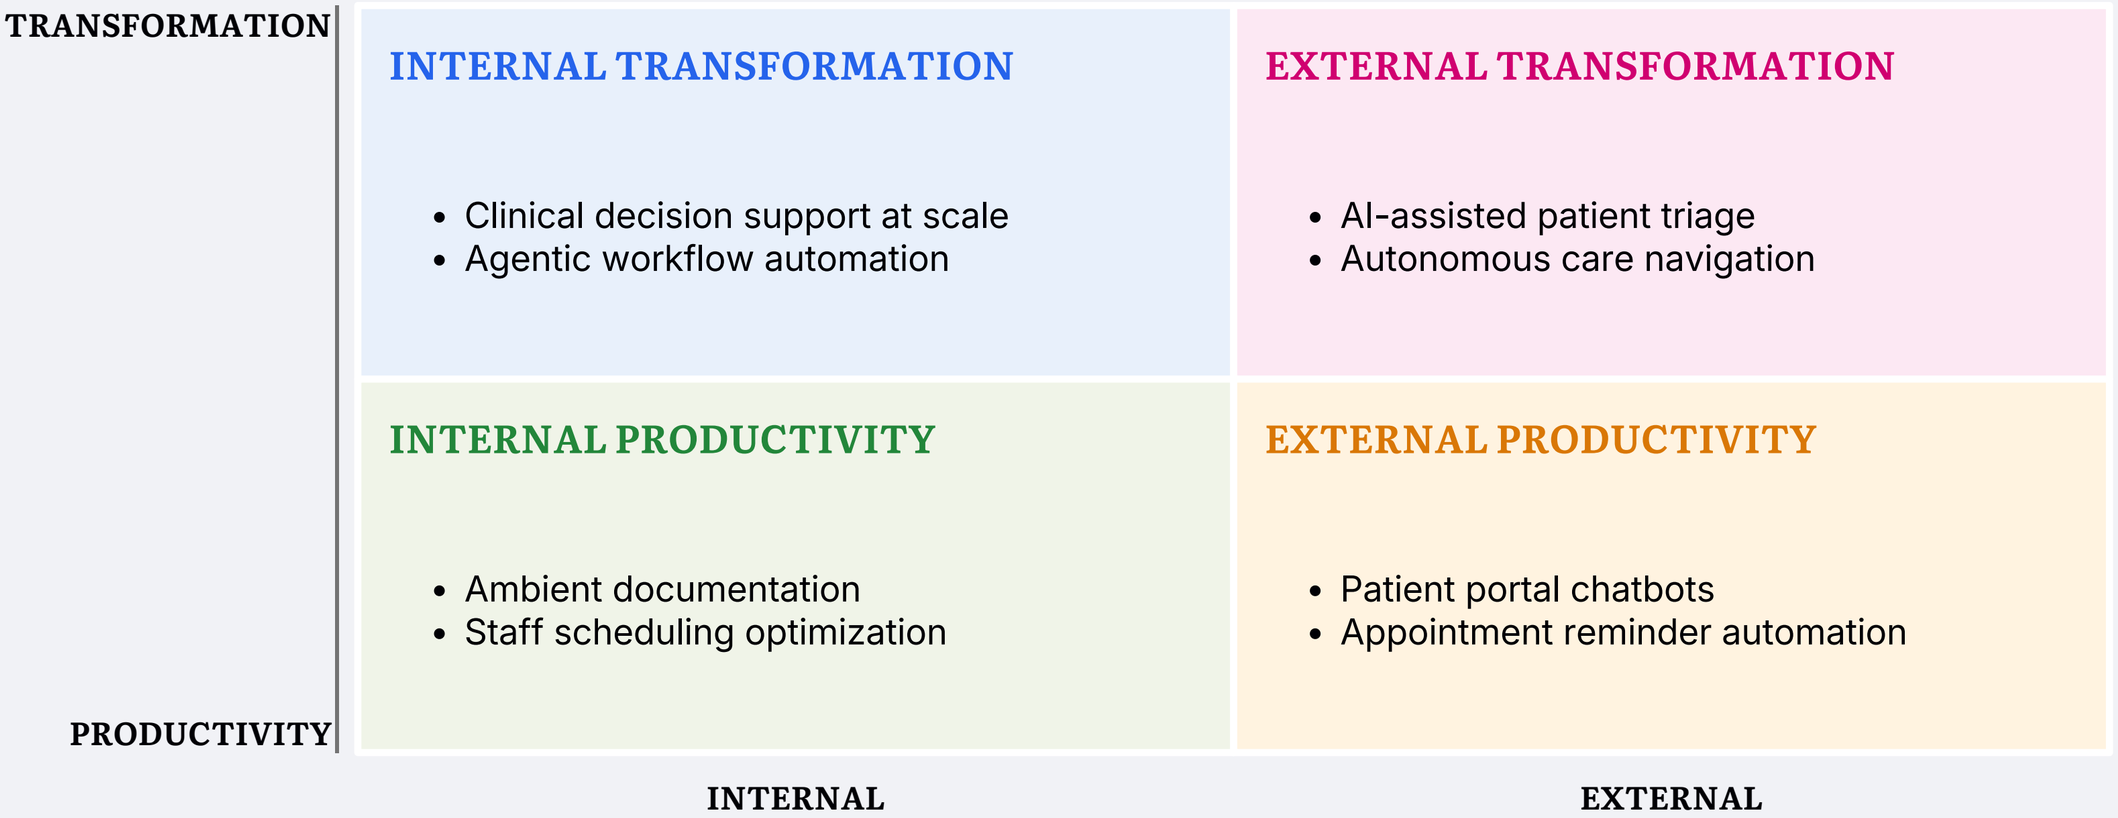

The goal is not zero risk. It's risk your organization can accept.

DOMAIN	INHERENT RISK	MITIGATIONS APPLIED	RESIDUAL RISK
Performance	HIGH	MITIGATIONS	LOW
Privacy	HIGH	MITIGATIONS	LOW
Security	HIGH	MITIGATIONS	MEDIUM
Ethical	MEDIUM	MITIGATIONS	LOW
Legal	MEDIUM	MITIGATIONS	LOW

Benefits are not generic. It's a matrix of who gains what.

TYPES OF BENEFIT

Financial
Revenue, cost savings, reimbursement capture

Operational
Cycle time, throughput, capacity unlocked

Clinical
Outcomes, quality of care, patient safety

Strategic
Market position, capability, competitive advantage

WHO BENEFITS

Patients
Access, experience, outcomes, safety

Clinicians
Cognitive load, time-to-decision, burnout

Administrators
Cost control, compliance, capacity

Payers & Partners
Risk sharing, coordination, value contracts

When risk and benefit sit on the same canvas, prioritization stops being political.

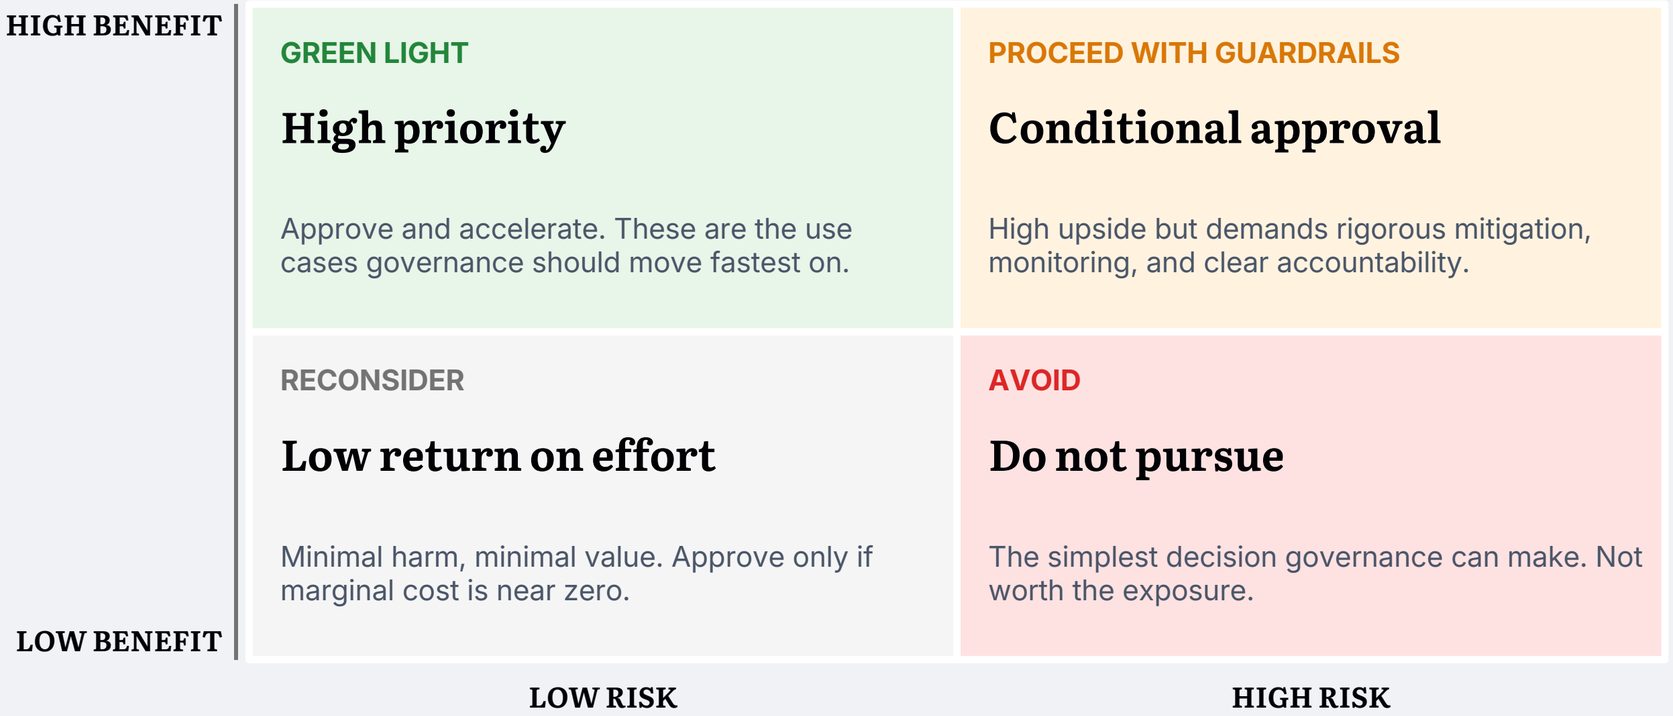

Discussion

Your CEO wants to know which AI initiatives are actually worth it. Your AI governance program should answer.

Scenario: Your organization has 35 AI use cases in flight or under review. The CFO has asked for a portfolio view: which ones are delivering measurable value, for which stakeholder groups, and whether that value justifies the risk and operational cost of each. No other team has a complete picture. Your governance program — through intake documentation, risk assessments, and benefit tracking — is the closest thing to a source of truth. What do you do with that?

Considerations

- *Does your current governance process capture expected benefits and stakeholder impact at intake, or only risk?*
- *How do you measure whether a deployed AI use case is delivering the value that was projected when it was approved?*
- *When risk and benefit are weighed against each other, who makes that call — and is there a consistent framework for doing it?*
- *How do you communicate differentiated impact across stakeholder groups: patients, clinicians, administrators, payers?*

Thank you!

Website	www.trustible.ai
E-mail	contact@trustible.ai
LinkedIn	https://www.linkedin.com/company/trustible/
HQ address	1201 Wilson Blvd, Floor 25, Arlington, VA 22209

Trustible