

SAMPLE DELIVERABLE · EXECUTIVE RISK BRIEFING

AI Agent Risk & Governance Clarity

Executive risk briefing and decision framework for AI agent adoption.

PREPARED FOR	Acme Corp
SECTOR	Enterprise Technology (Illustrative)
ENGAGEMENT	AI Agent Risk and Governance Clarity Workshop
SCOPE	Executive leadership team, AI agent adoption strategy
DELIVERED	April 2026

01 Executive Summary

Key findings from the AI Agent Risk and Governance Clarity briefing conducted with Acme Corp's executive leadership team.

CURRENT STATE

Acme Corp has **11 AI agents in production** across customer support, internal operations, and data analytics workflows. An additional 6 agents are in staging or pilot. The organization has strong IAM foundations (enterprise IdP, SCIM provisioning) but **no runtime governance layer** for agent behavior. Agents operate with static credentials, broad access, and no action-level policy evaluation.

Key Findings

AI agents are operating without runtime governance

FINDING

All 11 production agents are authenticated via IAM and hold valid credentials. None are subject to action-level policy evaluation. The gap between "the agent has access" and "the agent should take this action" is unaddressed.

Model safety is not agent governance

FINDING

The organization has invested in prompt guardrails and content filtering at the LLM layer. These controls address model behavior, not agent behavior. Agents choose tools, chain actions, and delegate work. These runtime decisions are not governed.

Agent inventory is incomplete

FINDING

The AI platform team is aware of 11 production agents. Discovery during the briefing identified 3 additional agents deployed by other teams operating outside the platform team's governance scope.

Headline Numbers

KNOWN PRODUCTION AGENTS 11 (+ 3 discovered)	AGENTS WITH RUNTIME GOVERNANCE 0 of 14
AGENTS WITH BROAD CREDENTIALS 14 of 14	DELEGATION CHAINS TRACKED None

02 AI Agent Exposure Assessment

Current-state snapshot of AI agent deployment across Acme Corp's enterprise, including systems accessed, credential posture, and governance gaps.

Agent Landscape

Workflow	Agents	Systems Accessed	Credential Type	Runtime Governance
Customer Support	3	CRM, Ticketing, Knowledge Base	Long-lived API keys	NONE
Data Analytics	4	Data Warehouse, BI Platform, External APIs	Service account tokens (90-day)	NONE
Internal Operations	4	HR System, Procurement, Scheduling	OAuth (user-delegated)	NONE
Shadow Agents	3	Various (discovered during briefing)	Mixed	UNKNOWN

Governance Gap Summary

ACTION-LEVEL POLICY EVALUATION

Not implemented

INTENT VALIDATION

Not implemented

DELEGATION CHAIN TRACKING

Not implemented

CREDENTIAL BROKERING

Not implemented

EXECUTION LINEAGE

Not implemented

HUMAN-IN-THE-LOOP ESCALATION

Ad hoc (not systematic)

KEY OBSERVATION

Acme Corp has strong identity foundations. The gap is not in who agents are, but in governing what they do. This is a common pattern: organizations invest in IAM for agents but lack the runtime governance layer that evaluates behavior in context.

03 Decision Framework and Next Steps

Recommended decisions and next steps for Acme Corp's executive team, sequenced by urgency and impact.

Immediate Decisions (This Month)

- DECISION 1

Complete the agent inventory. The 3 shadow agents discovered during this briefing indicate the inventory is incomplete. Before governance can be applied, the organization needs to know what is running. Assign ownership of agent inventory to the AI platform team with a 2-week deadline.
- DECISION 2

Rotate long-lived credentials. All 14 agents hold credentials with 90+ day lifetimes and broad scope. Begin immediate rotation and evaluate short-lived, scoped token issuance for the highest-risk agents (customer support, data analytics).

Near-Term Decisions (Next Quarter)

- DECISION 3

Conduct a threat model for the highest-risk agent workflows. Customer support agents access CRM data with customer PII. Data analytics agents access the data warehouse with broad read scope. A structured threat model will identify specific risks and recommended controls. *Recommended: AI Agent Security Threat Modeling Workshop.*
- DECISION 4

Evaluate runtime governance architecture. The gap between IAM and agent execution is the primary control gap. Evaluate whether to build, buy, or partner on the runtime governance layer. *Recommended: Authorization and Runtime Control Architecture Workshop.*

Strategic Decisions (Next 6 Months)

- DECISION 5

Establish an agent governance program. As agent adoption scales, governance must scale with it. Assign executive ownership, define the governance maturity target, and build the roadmap. *Recommended: Agent Governance Readiness Assessment.*

Recommended Workshop Path

NOW	NEXT	THEN	ONGOING
Risk Clarity	Threat Modeling	Architecture	Readiness
Acme Corp · Executive Risk Briefing - This briefing (completed) - Agent inventory - Credential rotation	- Map specific risks - Prioritize controls - Remediation plan	CONFIDENTIAL - Design control plane - Policy framework - Integration patterns	- Maturity assessment - Governance program - Continuous improvement