Covasant
AI-Driven. Human-Inspired

# The How-to of Governed Agentic AI

An executive field guide to deploying AI agents at scale, with the governance, control, and confidence that your organization demands.

---

## HFS Research identifies three bottlenecks preventing enterprises from scaling agentic AI: operating models, data readiness, and **governance maturity**.

*- HFS Research, Agentic Technology Report 2026*

**72%**
Are deploying agentic AI in their enterprises.
*- Insights from Mayfield's CXO Network 2026 Survey*

**137%**
rise in AI-related lawsuits in 2025 in US.
*- Forrester report of US AI Regulations*

**$1.5T**
Services-as-Software market by 2035
*- HFS Research*

**97%**
exploring agentic AI, only 12% have centralised governance
*- According to OutSystems' 2026 State of AI Development report*


**Covasant named by Google Cloud among top AI-native services partners.**

# Governance gaps hinder pilot to production scale

Agentic AI is no longer experimental. Enterprises are moving to production where the stakes are real. AI agents today are making decisions, interacting with customers, and accessing systems at machine speed. The critical question is how to govern them.

### Visibility Gap

Leadership doesn't know how many agents are deployed, what they do, or who owns them.

### Control Gap

When an agent misbehaves, there is no kill switch, no escalation path, no way to intervene.

### Accountability Gap

When something goes wrong, there is no audit trail and no one can answer 'what happened and why?'

## Regulatory Reality

### Compliance Gap

Different states in US have independent AI laws, creating a compliance patchwork.

### Rise in AI-related Lawsuits

In US there is 137% increase in AI-related lawsuits in 2025, and 32% reached class-action status.

### Reactive AI Governance

Companies must develop proactive AI governance programs that go beyond compliance.

## Six Steps to Enterprise-Grade Governed Agentic AI

01 **Discover & Register**

Build a unified Agent Registry so that every agent is known, owned, and classified by risk.

02 **Define Policies**

Codify acceptable actions, data boundaries, and escalation triggers before go-live.

03 **Human Oversight**

Design human-in-the-loop checkpoints and intervention controls into every workflow.

04 **Orchestrate**

A central orchestration layer prevents conflicting decisions across multi-agent workflows.

05 **Audit & Explain**

Every agent action logged, timestamped, and explainable, to regulators and boards.

06 **Measure & Scale**

Governance metrics reported at board level. Extend the framework to every new agent.

# AI Agent Control Tower
powered by CAMS

The Covasant AI Agent Control Tower, powered by CAMS, is a unified command center that provides real-time observability into agent performance & costs, ensuring explainability of every agentic decision across your enterprise.

### Agent Registry

Complete real-time inventory of every deployed agent, purpose, owner, permissions, version, and risk tier. No shadow agents. No surprises.

### Multi-Agent Orchestration

Coordinate fleets of specialized agents across complex workflows. Prevent conflicts, preserve context, and maintain a single control plane.

### Human-in-the-Loop Controls

Escalation paths and intervention controls built into every workflow from day one. Your CISO's kill switch. Your compliance team's audit anchor.

### Governance Policy Engine

Policy-as-code enforcement across your entire agent fleet. Real-time, context-aware, proportionate to risk, enforced before agents act.

## The Zero-Trust Architecture
## Built in Covasant AI Agent Control Tower

### Continuous Risk Management

Replaces one-time audits with real-time monitoring. Continuously evaluates the integrity of data, models, and agent behavior, providing ongoing assurance rather than point-in-time snapshots.

- Real-time behavioral monitoring
- Continuous data integrity checks
- Anomaly detection & alerting

### Explainable Outcomes

Every agent decision that blocks, modifies, or escalates must be understandable to both humans and systems. Critical for board-level trust, regulatory compliance, and effective governance at scale.

- Full decision audit trails
- Human-readable explanations
- Regulatory compliance ready

### Least Agency

It is the foundational security principle in agentic AI governance, places boundaries on what AI agents can decide and do. Agents receive the minimum permissions and authority needed to complete specific tasks, scoped by time and approval limits.

- Min. permissions per agent
- Time-bound access scopes
- Approval limits enforced

### AI Agent Control Tower Helps You Scale

Every agent, workflow and customer interaction is governed with visibility and control.

Goverance in Action

# **Industry Case Study:** Global Enterprise Asset Intelligence Leader

A global leader in enterprise asset intelligence, helping the world's largest retailers, manufacturers, and logistics providers track billions of items in real time, needed to govern their growing AI agent ecosystem with the same rigor they bring to physical asset visibility.

### The Challenge

- Agents deployed across partner onboarding, supply chain & field ops
- Shadow agents invisible to compliance teams
- Conflicting decisions across parallel agent workflows
- Audit failures: no traceable decision trail

### Our Approach

- Unified Agent Registry across entire ecosystem
- Multi-Agent Orchestration eliminating decision conflicts
- Human-in-the-loop for all compliance-critical decisions
- Real-time behavioral monitoring and full audit logging

### The Results

- 40% faster partner onboarding
- 100% audit compliance in Q1
- 60% reduction in agent conflicts
- 3x ROI on CAMS investment
- Board-level visibility achieved

## What Governed Agentic AI Delivers

### ISO/IEC 42001 Certified

The world's first international standard for an Artificial Intelligence Management System (AIMS)

### AI Agent Control Tower is a Part of CAMS

Covasant Agent Management Suite (CAMS) operationalizes global governance principles. By implementing the AI Agent Control Tower built on CAMS, enterprises leverage a platform architected to meet ISO/IEC 42001:2023 standards from the ground up.

CAMS governs every dimension of your agentic AI operation, from identity and data to security and compliance, so that you don't have to choose between moving fast and staying safe.

## Working With World's Leading Enterprises

ZEBRA · ESAB · macalvins · Olam Agri · NORTHERN TOOL + EQUIPMENT · NOKIA · motilal oswal · Unilever · BAT · vodafone · BMS

### About Covasant Technologies

Covasant Technologies is an emerging global leader in delivering Agentic AI-led services that address complex, industry-specific challenges. Through its pioneering Services-as-Software model, Covasant brings together capabilities in data engineering, digital and cloud, AI engineering, and Enterprise Risk Management (ERM). Strategically located in innovation hubs including Hyderabad, Plano, New York, Los Angeles, London, and Dubai, Covasant leverages a premier talent ecosystem to deliver scalable, autonomous solutions. Our "Governance-Led" approach ensures that global enterprises can automate decision-making and orchestrate business processes with the reliability and speed required in today's market.

Plano, USA • London, UK • Hyderabad, India • Dubai, UAE