



Sutherland · ServiceNow AI Control Tower

# From AI Sprawl to AI Command

How Sutherland helps enterprises turn  
ServiceNow AI Control Tower into a governed,  
value-generating asset



## EXECUTIVE SUMMARY

Every IT organization today is running two AI programs at once: the one it planned, and the one that grew on its own. Across banking, insurance, healthcare, manufacturing, retail, telecom, hi-tech, utilities, and travel, business units have embraced copilots, fraud agents, clinical assistants, personalization engines, network-optimization bots, and predictive-maintenance agents faster than IT has been able to inventory, secure, or govern them.

This is not a fringe problem. Enterprise AI tool spending is growing at roughly **40% year over year**, and an estimated 25% to 35% of that spend is occurring entirely outside IT's visibility. Separately, a **2026 study** of enterprise AI trends conducted by FTI Consulting for Smarsh found that while 55% of enterprises are actively deploying AI, only 26% believe their governance frameworks are keeping pace, and just 30% have the capability to detect AI running outside approved workflows.

The result is an AI estate that is delivering real value while operating largely outside the view of the people accountable for its risk. The specifics differ by industry: in banking and insurance it is a regulatory exposure, in healthcare a patient-safety and compliance risk, in manufacturing and utilities a safety and reliability issue, and in retail, telecom, hi-tech, and travel a direct hit to customer trust and brand. But the underlying challenge is identical everywhere. An ungoverned agent is rarely just a minor IT incident.

This paper examines how ServiceNow AI Control Tower gives IT and risk leaders a single, CMDB-linked platform to discover every AI asset, observe how it behaves at runtime, govern it against policy, secure it against misuse, and measure the value it returns, across any cloud, any vendor, and any business unit. It also makes the case, grounded in Sutherland's work with IT and risk leaders, that the platform alone is necessary but not sufficient. The organizations that get the most value pair AI Control Tower with a deliberate operating model: clear governance roles, industry-specific policy packs, and a continuous-improvement discipline that treats AI governance as a living capability rather than a one-time deployment.

# The Business Problem: An AI Estate Growing Faster Than the Ability to See It

Most enterprise AI did not arrive through a formal technology roadmap. It arrived the way shadow IT always has: one team at a time, one license at a time, justified by a real and immediate business need. A collections team adopts a conversational agent to triage overdue accounts. A claims unit stands up a model to flag suspicious activity. A plant manager deploys a predictive-maintenance bot from an OT vendor. Each decision is locally rational. Collectively, they produce an AI estate that central IT did not build, did not approve, and in many cases does not know exists.

The scale of that gap is no longer speculative. Active AI agents inside common productivity ecosystems have grown roughly 15 times year over year, and as much as 18 times in large enterprises, according to Microsoft's 2026 Work Trend Index. That growth rate is compounding a governance problem that was already unresolved for simpler, human-supervised tools. Separately, Deloitte's seventh annual State of AI in the Enterprise survey of more than 3,200 business and IT leaders across 24 countries found that only 21% of organizations currently have a mature governance model in place for agentic AI, even as roughly three-quarters expect to be using AI agents at least moderately within two years. Put simply, the deployment curve and the governance curve are moving in opposite directions, and the gap between them is the space in which incidents happen.



That gap also carries a direct cost. Organizations without established AI governance are estimated to spend roughly 2.5 times more on AI-incident remediation than organizations with controls already in place. That figure should reframe how leadership teams think about governance: not as a compliance tax on innovation, but as one of the more reliable cost levers available to an IT organization managing a fast-growing estate.

# Why the Pain Point Looks Different by Industry, and Why IT Owns All of It

The mechanics of AI sprawl are consistent across sectors, but the consequences are not, and understanding that difference matters for how leaders prioritize action.



## Banking, Financial Services, and Insurance

Fraud, underwriting, credit-decisioning, and conversational agents are often deployed by business units or fintech partners outside core IT's inventory. Model-risk regimes such as SR 11-7, RBI and MAS guidelines, DORA, and the EU AI Act require explainability and audit trails that IT must be able to produce on demand.



## Healthcare and Life Sciences

Clinical decision support, prior-authorization bots, and revenue-cycle agents touch protected health information across payer and provider systems. HIPAA exposure, algorithmic bias, and patient-safety risk turn an ungoverned agent into a clinical and legal liability, not just a ticket.



## Manufacturing and Utilities

Predictive-maintenance, quality-inspection, and grid-monitoring agents operate across converged OT and IT environments. A mis-governed agent on the plant floor or the grid can trigger safety incidents, production downtime, or a regulatory reporting failure, extending the blast radius well beyond the data center.



## Retail, Travel, and Hospitality

Personalization, dynamic-pricing, and loyalty agents are frequently deployed independently by e-commerce, merchandising, and marketing teams, or by third-party airline, hotel, and booking partners. Consumer-protection and data-privacy rules increasingly require explainability for decisions that touch customers directly and instantly.



## Telecom and Hi-Tech

Network-optimization, churn-prediction, and coding-assistant agents span multiple hyperscalers and fast release cycles. Lawful-intercept, data-privacy, IP-exposure, and service-reliability obligations mean an ungoverned agent is a regulatory and commercial exposure that can scale as quickly as the product itself.

What unites these industries is a simple fact: IT is accountable for the risk regardless of who deployed the agent. A regulator, a plaintiff's attorney, or a board does not distinguish between an agent that IT approved and one that a business unit stood up independently. The organization owns the outcome either way.

# The Implications: Operational, Financial, and Strategic

Operationally, an ungoverned AI estate means IT cannot answer basic questions with confidence: how many AI agents and models are running, what data they touch, who owns them, and what they are permitted to do. That uncertainty turns every audit, every security review, and every incident into an ad hoc scramble rather than a repeatable process.

Financially, the exposure compounds in more than one direction. There is the direct cost of remediation noted above, but there is also a slower-moving cost: value left uncaptured. **McKinsey's 2025 global AI survey** found that although the large majority of organizations now use AI regularly, only 39% report any enterprise-level EBIT impact from it, and a small minority qualify as high performers capturing outsized returns. Ungoverned, fragmented AI estates are a structural reason that value stays trapped at the use-case level instead of scaling to the enterprise. Without a shared inventory and a shared view of cost and performance, CIOs and CFOs cannot even agree on what is being spent, let alone what it is returning.

Strategically, the cost of inaction is regulatory as well as reputational. The **EU AI Act** sets administrative fines of up to €35 million or 7% of an organization's total worldwide annual turnover for the most serious violations, whichever is higher. That is not a distant, theoretical ceiling; it is the kind of number that changes how a board evaluates risk appetite for AI deployment generally, not just for the specific agent under scrutiny.





# What Prevents Organizations From Solving This

If the problem is this visible, why does it persist? In most enterprises, three things get in the way.

Governance has been treated as a policy document rather than a platform capability. A written AI policy is necessary, but it does not discover a single unauthorized agent, and it does not stop one from touching production data. The tools available have been point solutions bolted onto the environment after the fact: a dashboard here, a scanning tool there, none of them connected to the systems of record IT already trusts for infrastructure, identity, and workflow.

Most persistently, no single function owns the problem end to end. Discovery might sit with security, policy with risk, and remediation with IT operations, with no shared inventory connecting the three. The result is that even organizations that know they have a shadow AI problem struggle to act on that knowledge at the pace the problem is growing.

# A More Effective Approach: Governance as a Platform Capability

Sutherland's perspective is that IT organizations that treat AI governance as a platform capability, embedded in the same system of record that already manages their infrastructure, identities, and workflows, will out-execute those that treat it as a policy document or a disconnected point solution.

ServiceNow AI Control Tower is built on that premise. Running on the Now Platform, it enriches every AI agent and model as a configuration item, which means AI oversight becomes an extension of capabilities IT already trusts: the CMDB, identity and access management, and workflow automation, rather than another dashboard to reconcile.

**ServiceNow** has continued to expand the platform's reach: recent updates add dozens of new integrations spanning AWS, Google Cloud, and Microsoft Azure, along with enterprise applications such as SAP, Oracle, and Workday, and extend discovery to non-human identities and connected OT and IoT devices.

Sutherland's position is that the platform gives enterprises the technical control plane, but it does not run itself. The organizations that capture the most value are the ones that pair the platform with a deliberate operating model, industry-specific policy packs, and a continuous-improvement discipline.

The platform organizes its capabilities around five functions:

- **Discover** finds every AI asset across the enterprise, first-party or third-party, including agents, models, and MCP servers, closing blind spots before they become incidents.
- **Observe** provides continuous, runtime monitoring of agent behavior rather than periodic audits, giving teams visibility into how an agent is actually reasoning and acting.
- **Govern** applies AI-driven risk assessment and policy enforcement across agents, models, data sets, and prompts, not just generative tools.
- **Secure** enforces identity and least-privilege access on every AI asset, treating an agent's permissions with the same discipline applied to a human user's.
- **Measure** turns AI activity into cost and ROI dashboards that give CIOs and CFOs a shared, current view of what each AI investment is actually returning.

# From Strategy to Execution: How Sutherland Helps

Sutherland works with IT and risk leaders to move from AI sprawl to AI command in three phases, built around AI Control Tower's five native capabilities.

## PHASE 1



### **Advisory: Assess and Prioritize**

Sutherland runs an AI landscape and shadow-AI discovery assessment mapped to the client's CMDB, along with a gap analysis against the frameworks that matter for their sector: model-risk management for BFSI, HIPAA and clinical governance for healthcare, OT and IoT safety standards for manufacturing and utilities, consumer-data-privacy rules for retail and travel, subscriber-data and network obligations for telecom, and IP and release-velocity risk for hi-tech. The output is a prioritized roadmap and business case for AI Control Tower adoption.

## PHASE 2



### **Implementation: Configure and Integrate**

Sutherland delivers end-to-end configuration of all five AI Control Tower capabilities, including connectors into the major hyperscalers and enterprise systems noted above, identity and least-privilege enforcement for every agent, and industry-specific policy packs and approval workflows built directly into the Now Platform.

## PHASE 3



### **Continuous Improvement: Operate and Optimize**

Sutherland provides ongoing tuning of governance policies as new agents and models are onboarded, live ROI and value dashboards tied to business KPIs, and a co-managed operating model in which Sutherland's specialists monitor, triage, and continuously improve the AI estate so the client's IT team stays focused on strategic priorities rather than day-to-day policing.

What differentiates this approach is that each phase is customized to the industry it serves. A BFSI policy pack encodes model-risk and fair-lending requirements. A healthcare policy pack encodes HIPAA and clinical-safety review gates. A manufacturing or utilities policy pack extends governance into OT and connected-device identities. A retail or travel policy pack focuses on consumer-data consent and third-party agent oversight. A telecom policy pack addresses subscriber-data and network-reliability obligations, and a hi-tech policy pack is tuned to fast release cycles and IP protection. The same underlying platform, configured with the right guardrails for the right industry, is what turns a generic control tower into one that actually fits how a given business runs.

# What Becomes Possible: Illustrative Applications by Industry

The following are illustrative scenarios: realistic applications of AI Control Tower's capabilities mapped to the kinds of AI sprawl already common in each vertical, showing what becomes achievable once the platform and operating model are in place.





### **BFSI**

A regional bank running dozens of AI-enabled tools across fraud, collections, and customer service, many unknown to central IT, could use discovery and identity-governance capabilities to build a complete, CMDB-linked AI asset inventory and turn a weeks-long scramble for a regulator-ready audit trail into a same-day report.



### **Healthcare**

A multi-facility hospital network could bring prior-authorization and clinical-triage agents under a single governance layer, making every agent decision traceable to a policy and an owner, the kind of traceability that gives clinical leadership confidence to extend AI-assisted triage further.



### **Manufacturing**

A global industrial manufacturer could extend governance across IT and OT to surface predictive-maintenance agents running on plant-floor systems with no formal owner, using runtime observability and kill-switch controls to intervene before a misbehaving agent affects production.



### **Retail**

A multi-banner retailer could gain a single inventory of every pricing and recommendation agent deployed independently across e-commerce, merchandising, and store operations, with policies designed to catch a mistuned pricing model before it reaches customers during a peak sales event.



### **Telecom**

An operator running network-optimization and churn-prediction agents across multiple hyperscalers could close access gaps ahead of a regulatory data-privacy review by gaining visibility into every agent touching customer and network data.



### **Hi-Tech**

A software company whose engineering, sales, and support teams have each adopted their own AI coding assistants and customer-facing agents could establish a company-wide AI asset inventory and value dashboards, giving leadership its first accurate picture of AI ROI across business units.



### **Utilities**

A regional utility could govern grid-monitoring and outage-prediction agents running across converged OT and IT networks, with runtime observability designed to flag an agent operating outside its intended parameters during a severe-weather event, before it affects regulatory reliability reporting.



### **Travel and Hospitality**

A travel group with dynamic-pricing and customer-support agents spanning airline, hotel, and booking partners could gain a consolidated view of every third-party and first-party agent touching guest data, tightening access ahead of peak travel season.



## The Value at Stake

Taken together, a governed AI estate delivers value across six dimensions. Visibility replaces blind spots with a single, CMDB-linked inventory of every AI agent, model, and MCP server. Risk reduction follows directly from fewer ungoverned agents in production, which lowers exposure to the incident and remediation costs already discussed. Regulatory readiness replaces ad hoc, manual evidence-gathering with continuous, dashboard-ready reporting for model-risk, HIPAA, OT-safety, and AI Act-style audits.

Speed to value increases because business units can scale new AI use cases faster when governance is built in from the start rather than bolted on afterward. Cost and ROI transparency gives CIOs and CFOs a shared view of what each AI investment actually returns. And sustained operations become possible through a co-managed governance model that keeps policies current as the estate grows, without permanently expanding internal headcount.

The compounding effect of these six outcomes is a shift in IT's relationship with AI, from reactive cleanup to proactive command. That shift ultimately reaches the customer. A bank that can prove its fraud models are governed can approve legitimate transactions faster. A hospital network that trusts its triage agents can safely extend them to more patients. A manufacturer or utility that can see every agent across its plants and grid can turn predictive maintenance into a reliable promise rather than a pilot project. A retailer or travel brand that can govern its pricing and personalization agents can move faster without surprising its customers.

# What Leaders Should Do Next

The gap between AI deployment and AI governance will not close on its own. Every data point in this paper points toward that gap widening, not narrowing, over the next several quarters. Waiting for a visible incident before establishing a governance platform is, in practical terms, the most expensive path to the same destination.

Sutherland recommends that IT and risk leaders start with a focused, time-boxed AI landscape assessment: a discovery exercise that surfaces the AI agents and models already running across the enterprise, benchmarks them against the governance requirements specific to the organization's industry, and produces a prioritized roadmap for deploying and operating AI Control Tower. This is deliberately a small first step, designed to produce a concrete inventory and business case within weeks, not a multi-quarter commitment.



## Closing Perspective

AI agents are already making decisions inside banks, hospitals, factories, stores, networks, and utilities today, whether or not IT formally approved them. The question for enterprise leaders is no longer whether to govern AI, but whether that governance will be designed deliberately or discovered accidentally, after something has already gone wrong.

The enterprises that build a unified control plane now, one that discovers, observes, governs, secures, and measures AI as naturally as they manage every other technology asset, will be the ones able to scale AI with confidence while their peers are still explaining incidents to regulators, boards, and customers. The organizations that get there first will not only manage AI risk better. They will deliver a measurably better experience to the customers who depend on them.

**Contact Sutherland to schedule an AI landscape assessment workshop and see, in your own environment, how much of your AI estate is already outside your line of sight.**

Sutherland is an AI-driven business transformation partner to enterprises, including several Fortune 500 clients. We design, run, and automate business operations at scale, helping clients improve efficiency, strengthen customer experiences, and deliver measurable outcomes.

Most companies either build AI-led capabilities or run operations. Sutherland does both. We run complex, mission-critical processes while engineering the AI-led solutions that sit within them. This closed loop helps clients move beyond AI pilots to reliable, governed outcomes that improve and scale over time.

With deep operational expertise, advanced products, platforms, and engineering, we help clients modernize operations and build future-ready businesses.

