GOVERNING AI AT SCALE: THE AI CONTROL TOWER FLIGHT PATH (AICT)

Created by:

Nicholas Friedman

CEO & Founder of
Templar Shield

White Paper Title: Governing AI at Scale: The AI Control Tower FlightPath

Author: Nicholas Friedman - Templar Shield

Platform: Built on ServiceNow IRM, SecOps, and AI Control Tower Version: September 2025

Executive Summary

Artificial Intelligence (AI) is transforming how enterprises operate—from automating decisions to accelerating productivity. However, with rapid deployment comes a growing risk of ethical lapses, regulatory scrutiny, and shadow AI deployments. Enterprises need more than just point solutions. They require a unified, programmatic approach to govern AI.

Templar Shield's AI Control Tower FlightPath is a 6-phase enterprise roadmap designed to implement AI governance at scale—ensuring trustworthy, auditable, and regulation-aligned use of AI. Powered by ServiceNow's IRM, SecOps, and AI Control Tower capabilities, this roadmap operationalizes governance across policy, risk, controls, identity, and oversight—while accelerating innovation.

What is ServiceNow AI Control Tower

AI Control Tower is a centralized governance module on the ServiceNow AI Platform. It provides visibility, lifecycle management, risk & compliance controls, and performance monitoring for AI systems—both internal, third-party, and agentic.

It is built on ServiceNow's unified data architecture, leveraging the CMDB, Common Services Data Model (CSDM), configuration of AI assets (models, agents), and mapping them to business services.

Key capabilities include:

- AI Inventory / Discovery (agents, models, prompts, data sets)
- Risk & Compliance Management, including regulatory mapping (e.g. NIST AI RMF, EU AI Act)
- Automated workflows for AI use case review/approval, lifecycle changes, deployment / retirement of models.
- Performance & impact monitoring: continuously tracking metrics, usage, drift, anomalous behavior.
- Integration with IRM (Risk & Compliance), SecOps (for security incidents), case management, etc.
- Recent release updates (Zurich/May 2025) introduced agentic AI features, regulatory change management, and enhancements in IRM–AI risk integration.

Market Challenge: Fragmented AI, Unified Risk

- 42% of enterprises report having multiple AI deployments without centralized visibility.
- Global regulations such as the **EU AI Act**, **NIST AI RMF**, and **ISO/IEC 42001** mandate stricter accountability.
- Shadow AI and unapproved large language model (LLM) usage are introducing unmanaged legal, ethical, and security risks.

Without structured oversight, AI becomes a liability.

The Solution: AI Control Tower FlightPath

A **6-phase implementation roadmap** that combines:

- The six components of Templar Shield's AI Guardian solution
- ServiceNow's automation modules (IRM, SIR, CCM, AI Control Tower)
- A role-based rollout plan engaging Legal, Risk, Security, IAM, and Data teams

FlightPath Overview

⚔️ Templar Shield's AI Governance FlightPath™

A 6-Phase Enterprise AI Governance & Automation Roadmap on ServiceNow

This FlightPath is purpose-built to:

- Operationalize the 6 pillars of Templar Shield's AI Guardian solution
- Integrate ServiceNow's IRM, SecOps, and AI Control Tower into one unified governance program
- Implement trustworthy, auditable, and regulation-aligned AI at enterprise scale

🔶 PHASE 1: AI Strategy, Maturity & Footprint Discovery (Months 0–2)

Objectives:

- Identify shadow AI and known/unknown models
- Establish baseline AI governance maturity
- Map AI assets to business functions, data sources, and risk domains

🔧 Modules & IP Components

- AI Maturity & Footprint Assessment Accelerator
- ServiceNow CMDB, APM, and AI Inventory Tables
- AI Guardian Dashboards (Heatmaps, Scoring, Shadow AI Detection)

🧠 Key Outcomes

- Discover all current AI/ML models, APIs, datasets, and LLM usage
- Identify shadow AI and unmanaged deployments
- Establish baseline AI maturity score across 5 domains

🔁 ServiceNow Automation

- Populate CMDB CI Models for AI assets
- Enable AI Inventory Dashboard and risk profile scoring
- Link AI assets to business services and data sources

Participating Teams:

- Enterprise Architecture (EA) – lead model and system discovery
- IT Operations / CMDB Team – configure CI model classes
- Data Governance Office – help identify datasets feeding models
- Risk & Compliance – assess governance maturity baseline

🔶 PHASE 2: AI Policy, Risk, and Control Foundation (Months 2–4)

Objectives:

- Define AI-specific policy framework (aligned to EU AI Act, NIST AI RMF, etc.)
- Establish AI risk taxonomy and link to IRM modules
- Map regulatory content to actionable controls

🔧 Modules & IP Components

- Templar Shield's Regulatory Content Pack (EU AI Act, NIST AI RMF, ISO/IEC 42001)
- ServiceNow IRM: Policy Management, Risk Register, Control Library
- UCF Integration for continuous compliance updates

🧠 Key Outcomes

- Align AI operations to 20+ AI laws and standards
- Author and deploy enterprise-wide AI policy frameworks
- Establish tiered risk levels (high-risk, limited-risk AI) per EU AI Act

🔁 ServiceNow Automation

- Map AI risks to use cases and controls in real-time
- Enable AI-specific regulatory change tracking (RCM)
- Auto-generate policy exceptions and control test plans for AI systems

Participating Teams:

- Legal/Compliance – own policy authorship and oversight mapping
- Risk Management – define AI risk types, scoring, and tiering logic
- Audit & Internal Controls – review control effectiveness criteria
- IT Security – input on technical control mapping

🔶 PHASE 3: AI Control Testing, CCM, and Real-Time Monitoring (Months 4–7)

Objectives:

- Deploy automated control tests (CCM) for AI systems
- Enable real-time dashboards for policy enforcement
- Connect failed controls to issue/risk workflows

🔧 Modules & IP Components

- ServiceNow CCM, SPC, and Agentic AI + Now Assist
- Prebuilt Automated Controls from Templar Shield
- AI-specific compliance dashboards, control health scoring

🧠 Key Outcomes

- Monitor AI controls continuously: fairness, explainability, data integrity
- Trigger automatic remediation and issue logs when AI deviates
- Strengthen oversight without requiring constant manual audits

🔁 ServiceNow Automation

- CCM auto-tests AI control efficacy daily or on model update
- Failures trigger issue logs and risk score adjustments
- Live dashboards show pass/fail rates by AI system or risk domain

Participating Teams:

- Compliance Operations – define control logic, thresholds, and schedules
- IRM Engineering / Platform Team – build automated control tests
- Cybersecurity/IT Security – monitor AI security posture & anomalies
- ServiceNow Admins – configure dashboards, workflows, data sources

🔶 PHASE 4: AI Risk Management & Incident Response (Months 7–9)

Objectives:

- Classify and triage AI-specific incidents (bias, drift, hallucination)
- Link AI incidents to model metadata and known risks
- Enable root cause analysis and forensic capability

🔧 Modules & IP Components

- ServiceNow Risk Management, Security Incident Response (SIR)
- ServiceNow AI Control Tower: for centralized oversight
- IRM-VR integration for proactive AI risk detection

🧠 Key Outcomes

- Track vulnerabilities from biased data, adversarial prompts, drift
- Enable real-time response to AI failures or regulatory violations
- Link AI incidents to models, datasets, and downstream decisions

🔁 ServiceNow Automation

- Auto-create risk entries when AI model metrics breach thresholds
- Auto-escalate incidents to compliance or security depending on type
- Use MITRE ATT&CK-like AI threat libraries for root cause classification

Participating Teams:

- SecOps / SOC Team – triage AI-related security and incident tickets
- Risk & Governance – manage AI risk response workflows
- Platform Ops – enable cross-module linkage to CMDB and AI metadata
- Legal / Ethics Office – review incidents with legal/regulatory exposure

🔶 PHASE 5: Identity, Access, and Prompt Governance (Months 9–11)

Objectives:

- Implement role-based access to AI models and prompts
- Prevent prompt injection, unauthorized LLM use, and SoD violations
- Use prompt libraries for compliance, audit, risk teams

🔧 Modules & IP Components

- Clear Skye IGA (or SailPoint), IRM Access Review, Data Loss Prevention (DLP)
- Industry-Specific Prompt Libraries (for compliance, audit, security use)
- Role-based prompt access in ServiceNow Workspaces

🧠 Key Outcomes

- Prevent over-provisioning to models or prompt injection attacks
- Govern access to AI-generated decisions, training data, and model APIs
- Enforce human review, training completion, and SoD before granting prompt access

🔄 ServiceNow Automation

- AI Agents and systems treated as identities (CI class)
- Approval rules: training complete → role assigned → prompt unlocked
- Log prompt usage, analyze for risk categories (e.g., compliance, security)

Participating Teams:

- Identity & Access Management (IAM) – configure JML, SoD, approval flows
- Internal Audit & Compliance – approve use-case-based prompt access
- Training & HR – enforce training prerequisites for AI prompt users
- Platform & Workspace Teams – manage UI roles, tracking, and dashboards

🔶 PHASE 6: Unified Oversight, Reporting & AI Governance COE (Months 12–14)

Objectives:

- Launch full AI Governance dashboard (AI Control Tower)
- Operationalize reporting to regulators, boards, and internal committees
- Transition ownership to AI Governance Center of Excellence (COE)

🔧 Modules & IP Components

- ServiceNow AI Control Tower, IRM Dashboards, PPM, Performance Analytics
- AI Governance Workspace (central homepage for AI oversight)
- Quarterly attestation templates, risk/ethics dashboards

🧠 Key Outcomes

- Real-time governance dashboards by model, business unit, regulation
- Quarterly AI oversight reports for boards, regulators, auditors
- Launch of permanent AI Governance Center of Excellence (COE)

🔄 ServiceNow Automation

- Auto-generate compliance reports by jurisdiction (e.g., EU AI Act Annex III)
- Dashboard integrations show AI adoption KPIs vs governance health
- Embed risk and model metadata into PPM for project-level oversight

Participating Teams:

- AI Governance Office / COE – program ownership and policy authority
- Board Risk Committees – receive quarterly oversight reports
- Business Units – maintain model-level compliance & operational guardrails
- IRM, SecOps, ITSM, APM Teams – sustain platform integrations and automations

Best Practices for Implementation & Integration

Based on what ServiceNow publishes and what practitioners report, here's a best-practices approach for implementing AI Control Tower well and integrating it with IRM, SecOps, and other modules:

Practice	Why It Matters	Key Steps / Recommendations
Start with Inventory & Discovery	You can't govern what you can't see. Shadow AI, dispersed models, agentic workflows often exist without oversight.	• Use AI Control Tower's asset discovery to catalog all models, agents, prompts, datasets. • Ensure integration or ingestion from third-party sources, bespoke models, legacy deployments. • Populate the CMDB/CSDM properly, with metadata: owner, business function, data sensitivity, deployment stage.
Define Risk and Policy Frameworks Early	To provide consistent criteria for evaluation, classification, approval, audit. Helps avoid subjective or fragmented governance.	• Align with relevant regulatory frameworks (e.g. EU AI Act, NIST AI RMF). • Define risk tiers: data sensitivity, decision automation vs advisory, domain (finance, health, etc.). • Develop policy templates (privacy, fairness, explainability) that map to controls in IRM.
Implement Lifecycle & Review Workflows	AI models evolve. Changes (data, model, prompt) often introduce new risks. Lifecycle governance ensures upstream reviews and downstream monitoring.	• Use AI Control Tower to incorporate workflows for proposal → review → approval → deployment. • Trigger reviews for model updates, data changes, drift detection. • Include stakeholders across legal, risk, security, business units.
Integrate with IRM and SecOps Modules	Governance, risk, compliance, and security are deeply intertwined. Integration avoids gaps and overlaps.	• Map AI risk and compliance with IRM: risk registers, policy library, control objectives. • Link security incidents or vulnerabilities from SecOps (e.g. SecOps incident response, vulnerability response) that involve AI models to AI Control Tower's risk tracking. • Use continuous control monitoring (CCM) to test AI-related controls (data privacy, model robustness, input/output validation).
Enable Real-Time Monitoring & Alerting	AI can drift, be misused, or produce harmful outputs quickly.	• Collect metrics (Performance, fairness, bais, error rates, usage volume) and establish thresholds. • Automate drift detection and anomalous behaviour using AI Tower + monitoring modules.

	Real-time insight helps detect and respond early.	• Set up alerts, notification workflows, and SLAs for response.
Ensure Transparent Human Oversight & Auditable Trails	Regulatory, ethical, and stakeholder trust demands that decisions can be traced,human override exists, and audit records are preserved	• Maintain metadata: who initiated model, who approved, what data was used, version used, prompt or configuration. • Provide tools for explainability (why did the model make this decision) where applicable. • Enable human in the loop / override points for high-risk decisions.
Govern Access & Identity for AI Assets	Because AI agents and models can access sensitive data, perform automated actions, or impact business-critical outcomes.	• Use IAM / IGA modules to treat models/agents as identities or CIs where appropriate. • Define role-based access controls for model deployment, prompt usage, configuration, approvals. • Conduct regular attestation / review of who has access, and track prompt / use-case histories.
Align with Business Strategy & Outcomes	Governance must map to value; otherwise, it becomes a compliance burden. Prioritize initiatives that deliver business impact.	• Tag AI initiatives with business service alignment, KPIs, value metrics. • Use dashboards to show ROI, risk exposure, cost-benefit. • Prioritize resources on high-impact, high-risk AI use cases.
Scalable, Modular Deployment	Not all AI functions are equal; rolling out governance should accommodate varied risk and scale gradually.	• Start with pilot domains or high-risk models first, then expand. • Use configurable control templates. • Build modular integrations: IRM, CCM, SecOps, AI Control Tower so they can be adopted incrementally.
Maintain Regulatory & Ethical Agility	Laws and public norms are evolving. Global requirements (EU, US, APAC) will continue to shift.	• Use ServiceNow's regulatory change management and policy library to stay updated. • Update control definitions, workflows, risk thresholds as regulations evolve. • Incorporate ethical advisory or review board feedback.

Governance by CrossFunctional Teams	Different disciplines—business units, data, legal, security, risk—see different risks and must be aligned.	• Establish an AI Center of Excellence or committee. • Define clear ownership: who owns policy, risk scoring, model approvals, performance metrics. • Include business unit owners in lifecycle reviews.
Training, Change Management & Culture	Technology alone is not enough. Staff must understand why governance matters and how to comply.	• Provide training on AI ethics, risk, bias, explainability. • Use “guardrails” and accessible guidance for model developers, citizen developers. • Offer templates or playbooks for use cases, prompt design, etc.

Components from Templar Shield’s AI Guardian

Component	FlightPath Phase	Capability Area
Regulatory Content Pack	Phase 2	Policy, Risk, Compliance
Agentic AI Training Program	Phase 5	Access, Education
AI Maturity Assessment Accelerator	Phase 1	Discovery, Benchmarking
Prompt Libraries	Phase 5	Workspace Security
Auditable ServiceNow UI	Phase 5–6	Tracking & Transparency
AI Control Tower & IRM	Phase 3–6	Risk, Controls, Oversight

Suggested Integration Points & Architecture

Here are recommended integrations and architectural components when implementing AI Control Tower in a ServiceNow ecosystem:

Integrated Module	Purpose / Use Case	Key Integration Points
IRM (Governance, Risk & Compliance)	To manage policy, risk, and compliance artifacts for AI systems	- Map AI models to policies, controls, risk registers - Use IRM for impact assessments, compliance attestations, regulatory mapping - Use Regulatory Change Management (RCM) for AI-specific regulation updates
CMDB / CSDM	Maintain asset & model inventory, relationships	- Define CI types for models, agents, data sets - Maintain dependency graphs (models → data → deployments → business services)
Continuous Controls Monitoring (CCM) / SPC	For automated control testing of AI governance policies	- Build control checks for data quality, bias, model drift, prompt safety - Feed results back into risk scoring in AI Control Tower / IRM
Security Incident Response (SIR) & Vulnerability Response (VR)	To capture incidents or vulnerabilities arising from AI model operations or adversarial vectors	- Link AI model issues (e.g. misuse, data leakage) to SIR case workflows - Vulnerabilities in model libraries or dependencies enter VR process - Incident dashboards surfaced in AI Control Tower for executive oversight
Identity & Access Governance (IGA / IAM)	To control who can build, deploy, prompt, or alter AI systems	- Treat models/agents or their management interface as identities or CIs - Use role-based approvals, SoD, prompt libraries access - Perform periodic access reviews and attestation
Now Assist / Agentic AI	To automate governance tasks (summaries, mapping, control generation, issue resolution)	- Use Now Assist to reduce manual burden - For example, summarizing impact assessments, identifying redundant controls, generating draft policies or control objectives
Policy & Regulatory Change Management	To stay aligned with changing laws, ethics, and standards	- Maintain mapping from regulatory frameworks to internal policies/controls - Trigger reviews when regulations evolve - Use AI Control Tower dashboards to show compliance posture relative to regulation

Key Outcomes

- **30%+ increase in AI asset visibility**
- **80% reduction in audit preparation time** via automation
- **25% improvement in model compliance scores** (EU AI Act/NIST)
- **Real-time dashboards** for control health, prompt usage, and AI asset inventory
- **Full traceability** between model input/output, policy, control, and oversight

Use Case Examples

Use Case	Enabled By
Discovering Shadow AI	CMDB, APM, Maturity Assessment
Flagging High-Risk Prompts	Prompt Library, Access Review, CCM
Linking AI Drift to Incident	AI Control Tower + SIR + Risk Register
Regulatory Mapping (EU AI Act)	Regulatory Content Pack + IRM
Quarterly Board Reporting	AI Control Tower + IRM Dashboards

Common Pitfalls & How to Avoid Them

Pitfall	Risk	Mitigation / Lessons Learned
Shadow / Unmanaged AI Assets	Untracked models or agents introduce risk (data breaches, regulatory noncompliance)	Invest heavily in discovery phases; require registration of AI use cases; integrate across all cloud/data environments
Overcentralizing & Slowing Innovation	If governance is too onerous, teams may bypass it or delay deployment	Use tiered risk approach so low-risk AI has lighter controls; build guardrails rather than blocks; define SLAs for governance reviews
Data Silos & Poor Metadata	Difficulties tracing lineage, understanding data sensitivity, correlating usage to risk/exposure	Enforce consistent metadata standards; ensure CMDB/CSDM completeness; automate connectivity with data catalogs
Lack of Clear Ownership	Governance suffers if roles (legal, compliance, AI engineers, business units) are not clearly defined	Establish an AI CoE or governance committee; define roles for policy, risk, performance, security; assign executive sponsor
Neglecting Human Oversight / Explainability	Models deliver output that's unexplainable or decisions not attributable; risk to trust and compliance	Implement visible control mechanisms; document decisions; provide override paths; include explainability in requirements
Not Updating Controls & Risk Frameworks	Regulatory & ethical environments evolve; model behavior or usage drifts over time	Use regulatory change management; commit to regular reviews of risk taxonomies and policy updates; monitor model drift; retire outdated models
Insufficient Monitoring & Alerts	Issues (bias drift, errors, misuse) may go unnoticed until damage occurs	Define key metrics; set thresholds and alerts; automate anomaly detection; ensure visibility in dashboards

Success Factors

1. **Cross-Functional Governance** – Engage Legal, Risk, Data, Security, IAM, and Business Ops
2. **Automation First** – Use ServiceNow CCM, IRM, and AI Control Tower to reduce manual workload
3. **Continuous Improvement** – Reassess AI maturity and retrain users quarterly
4. **Compliance by Design** – Align all AI development and access to regulatory and ethical controls

Why Templar Shield

- Elite ServiceNow Partner with domain expertise in IRM, SecOps, AI Governance
- Creators of the **AI Guardian** solution suite (Reg Content Pack, Prompt Libraries, Maturity Frameworks)
- Proven delivery across highly regulated industries: BFSI, Healthcare, Critical Infrastructure, Public Sector
- Author of Governing Giants – Enterp[rise Risk Management's Role in AI Governance.

AI is the future—but only if it's governed. The AI Control Tower FlightPath is your launchpad to scale AI governance with confidence, compliance, and control.

Want to learn more or talk to our team of experts?
Contact us at info@templarshield.com or visit us at www.templarshield.com