HCLTech | Supercharging Progress™

Many Agents, **One Control** Plane

Placing AI Governance at the Heart of Service Management



Overview

In recent months, we have all seen the pace at which AI has been growing – both in terms of capability and in terms of economic impact.

Recently Forbes published 'The Next AI Governance Challenge: Governing Agent Behavior' showing the direction industry and governing bodies are taking : clear signal that as Agentic AI is moving from pilots to production at an industrial scale, bringing with it a shift in the multi-vendor landscape that enterprises know needs to be well governed to maintain control and generate the right ROI.

It's no longer just about many vendors, but also about each of them shipping multiple Agentic AIs—specialized for IT, HR, customer support, supply chain optimization and more — which enterprises can stitch together into an agentic mesh.

This means significant capability improvements, but also new points of failure, since uncoordinated agents can lead to duplicate calls, conflicting decisions, runaway token costs and fractured audit trails.

For business leaders, the question is: when many agents act, what guarantees coherence?

The solution is AI governance embedded into Service Integration and Management (SIAM), creating a single plane of control that enforces accountability, coordination and cost visibility across vendors and agents.

Are you prepared for this evolutionary leap?



The Maturity Curve: From Experimentation to Federated Ecosystems

While AI has taken over the strategic discourse across industries, safe and risk-free adoption remains a challenge. To succeed, enterprises need a very different kind of maturity and that evolution can be defined as three broad levels:

1. **Experimentation** —Isolated pilots and proof-of-concepts where AI assists within narrow functions.
2. **Operationalization** — AI embedded into workflows and service processes, but still siloed by vendor or business unit.
3. **Federated agent ecosystem** — A state where multiple agents from multiple vendors collaborate autonomously across enterprise and partner environments, guided by unified governance and shared observability.

As organizations move toward this third stage, the line between service integration and AI integration is disappearing. Coordinating systems, vendors and autonomous agents now demand a single, unified integration capability. This convergence redefines what operational maturity looks like, with the focus shifting from just reliable service delivery to real-time orchestration.

However, most enterprises are far from this ideal. AI adoption has raced ahead of governance readiness. A joint HCLTech–MIT Technology Review study found that:

- **87%** of executives see Responsible AI as critical
- **85%** admit they're not prepared to implement it at scale

Bridging the gap is essential for organizations looking to thrive in the new agentic mesh reality.

Emerging Governance Challenges



Imagine a retailer using pricing and inventory agents supplied by different vendors. In some scenarios, these agents could make simultaneous API requests to the same external LLM, causing a spike in token consumption and an API quota outage that results in stalled pricing updates.

This example highlights how **Agentic AIs don't fail in isolation**. While traditional automation followed scripts, Agentic AI pursues goals independently, developing adaptive protocols for inter-agent communication, which emerge without human design or oversight. Beyond higher costs and slower incident response, this opacity can create gaps in accountability and weak decision audit trails.

Multi-tenancy is another critical factor. Some agents are deployed as shared services across multiple organizations that might be direct competitors. In such cases, aggregated signals or poorly isolated context can leak commercial insights across tenants, multiplying competitive-intelligence risks.

As agentic deployments scale, addressing these concerns with effective governance will require a shift from checklist-style controls to operational, continuously calibrated systems. This is where **AI governance and service management converge.**

How SIAM Serves As Your Control Plane for Agentic AI

The governance challenges of the agentic mesh cannot be solved by policy alone. They demand a real-time approach to coordinating actions across vendors and agents. But while tackling AI governance in isolation is complicated, building on existing capabilities can be far more practical and significantly faster.

Recent research from Gartner establishes a strong connection between governance maturity and sustainable AI outcomes. The study found that **45% of organizations with high AI maturity** have kept their AI projects operational for three years or more, crediting robust governance and engineering practices as key enablers.

This need for a robust governance, combined with inherent advantages of using an existing service integration frameworks such as SIAM which already has linkages and understanding of organizational needs creates an undeniable logical argument, rather than building new control structures from scratch.

At its core, SIAM already provides the fundamental governance needs, including a single point of accountability, a framework for integration and end-to-end visibility across providers. What it did for human-managed services, it must now do for autonomous, intelligent ones.

The parallels are clear:

Agent registry → Service catalog / CMDB
Allows enterprises to maintain a single source of truth for every agent, including its owner, data scope and retrain cadence, so that accountability is built into daily operations.

Data-access and consumption policies → Change and access management
Ensures that every model or data request follows defined quotas, approval paths and security controls.

Orchestration layer → Event and incident management
Sequences agent actions, resolves conflicts and prevents cascading failures across services.

Observability → SLA and performance management
Brings together decision lineage, performance metrics and token consumption to monitor compliance and cost efficiency.

Vendor and contract governance → Contract and relationship management
Embeds AI-specific clauses on data provenance, audit rights, tenancy safeguards and accountability into supplier frameworks.

A Control Plane for Agentic AI

AI Lifecycle View	Creation of an AI lifecycle view covering: **Inception (reasons and triggers) → Contextualization → Implementation → Maturity → Currency and evolution → Retirement**
Guardrails for AI Implementation	Needs to govern the applicability of AI: • Each AI must have a purpose and aim specific to the organizational needs. • Well-defined and authorized operational boundaries — such as tech stack, data exposure, data consumption and boundaries including legislative obligations and security considerations.
AI System of Records and Engagements	• **AI agent definition** — each agent's composition: persona, memory, tools and model (such as LLMs). • **AI agent identity** — nomenclature, type, application. • **Consumption data structures.** • **Topology in the ecosystem** — structures and interaction points with systems and applications. • **Ontology** — inter-agent interactions, dependencies, hierarchies.
AI Agents Ownership and Accountability	Although AI agents can take autonomous decisions, the accountability must lie with the owner that deploys, maintains or utilizes the AI's capabilities. A central governance ensures each AI has an accountable owner.

In essence, these capabilities form the **SIAM-led control plane**, turning AI governance from a policy layer into an operational discipline where oversight, orchestration and optimization happen continuously, not retroactively.

Why This Matters Now

When AI governance is embedded into SIAM, control stops being a constraint and becomes an enabler, giving enterprises the clear visibility and confidence to scale rapidly and responsibly.

The choice is simple:

Accept **fragmented, ad-hoc integration** that drives up costs and obscures accountability

— or —

Adopt **a control plane** that transforms the agentic mesh from a source of chaos into a source of competitive strength.

Today, AI proliferation is an operational reality. In this multi-vendor, multi-agent landscape, mastery will come not from the models you run, but from the governance that keeps them aligned.

HCLTech's SIAM Governance Control Framework serves exactly this need : allowing you to be in control and ensuring the Agentic Ais deployed deliver the right value within the right framework, while generating the ROI promised to the business.



About the Authors



Arvind Raman

Global Head of Unified Service Management Business,
HCLTech

Arvind Raman is Senior Vice President and Global Head of Unified Service Management at HCLTech, bringing over 24 years of experience in enterprise service management, digital transformation, IT operations and platform-led innovation. He leads the strategy, growth and delivery of HCLTech's Unified Service Management portfolio, helping global enterprises modernize service delivery, optimize operations and accelerate business transformation through AI, automation and intelligent workflows.

With deep expertise in ServiceNow, SIAM, AI-powered operations and enterprise governance, Arvind partners with organizations to build resilient, experience-centric operating models that deliver operational excellence, business agility and measurable outcomes. He is a strong advocate for harnessing Agentic AI, unified workflows and enterprise governance to enable secure, scalable and responsible AI adoption.

A recognized thought leader and trusted advisor, Arvind regularly shares insights on AI governance, autonomous enterprises, ServiceNow innovation and the future of enterprise service management. He collaborates closely with customers and strategic technology partners to help organizations move from AI experimentation to enterprise-wide transformation with confidence.



Prashant Bhat

Practice Director,
HCLTech

Prashant Bhat is practice director for SIAM and USM at HCLTech. In this role, he collaborates closely with the Product Management Group (PMG) and Centers of Excellence (CoEs) to drive next-generation IT Ops and Service Management solutions for both existing and prospective clients, providing strategic advisory and pre-sales leadership.

He has played a pivotal role in building and nurturing HCLTech's SIAM capabilities since inception, which has helped maintain the long-standing, multi-decade thought leadership position for HCLTech. Through his deep industry expertise and client-centric approach, he has helped strengthen HCLTech's position as a trusted digital transformation partner in the financial services sector.

Prashant's charter is focused on identifying and developing new-age solutions, identifying new business opportunities, expanding existing client engagements and driving revenue growth within the USM domain. He is responsible for shaping the unit's business strategies, creating market momentum and enabling sustainable growth by delivering innovative, AI-enabled service management solutions that address evolving client needs.

HCLTech | Supercharging Progress™

hcltech.com