


WHITE PAPER · GCC ENTERPRISE AI GOVERNANCE

AI Control Tower: Build vs Buy and the Governance Architecture Decision

A synthesis of evidence for enterprise leaders

Synthesising ServiceNow's Knowledge 2026 announcements, McKinsey QuantumBlack's enterprise agentic platform architecture framework, and GCC regulatory developments for enterprise leaders navigating the AI governance decision in 2026.



AI Control Tower: Build vs Buy

C O N T E N T S

- 01 Executive Summary

- 02 The Problem: Deployment Without Defensibility

- 03 GCC Regulatory Considerations: PDPL, QCB, Dubai AI Seal, and EU AI Act

- 04 What Defensible AI Looks Like in Practice


- 05 What ServiceNow AI Control Tower is in 2026

- 06 The McKinsey Architecture Framework: Composable and Compostable

- 07 The Competitive Landscape: Choosing the Right Governance Architecture

- 08 The Vertex Diagnostic and Five Outcome Lanes

- 09 Avero's Point of View and Recommended Actions



70%

OF ORGANISATIONS FAIL TO REALISE FULL SERVICENOW PLATFORM VALUE.

AVERO WAS BUILT TO CLOSE THAT GAP

Apty / ServiceNow Digital Adoption Analysis, 2024-2025

ABOUT

About Avero

Avero is a ServiceNow-pure-play advisory and delivery firm operating across GCC sovereign and financial institutions. Built around a single platform, a focused delivery model, and a governance-first approach. The Vertex Framework™ is Avero's proprietary architecture decision methodology for GCC enterprise AI programmes.

About This Paper

Sections 02 through 04 establish the evidence base: the problem, GCC regulatory mandates and what practical defensible AI is. Sections 05 to 07 provide the platform, the architecture framework, and the competitive landscape. Section 08 introduces the Vertex Diagnostic and Five Outcome Lanes - Avero's proprietary framework. Section 09 sets out Avero's advisory position and three recommended actions.

Contact

contact@avero.works · <https://avero.works>

Authored by Avero. May 2026.

SOURCES

- QuantumBlack, AI by McKinsey, *Creating a Future-Proof Enterprise Agentic Platform Architecture* · April 2026
- ServiceNow Knowledge 2026, AI Control Tower, Action Fabric, AI Gateway, Microsoft Agent 365 · May 2026
- McKinsey, *The State of AI in 2026* · 2026
- McKinsey, *State of AI Trust in 2026: Shifting to the Agentic Era* · 2026
- Mordor Intelligence GCC Digital Transformation Report · 2025
- World Bank GCC Report · 2025
- GSMA MENA Report · 2025
- PDPL Enforcement Guidance · Saudi Arabia · 2025-2026
- QCB AI Governance Framework · Qatar Central Bank · 2025-2026
- Dubai AI Seal Programme Documentation · Dubai Centre for Artificial Intelligence (DCAI) · 2025-2026
- Dubai Department of Finance AI Procurement Directive · October 2025
- EU AI Act Enforcement Guidance · European Commission · 2024-2026
- Gartner CIO Survey · 2025
- Grant Thornton 2026 AI Impact Survey · April 2026
- Deloitte State of AI · March 2026
- Forrester 2026 Technology & Security Predictions · October 2025
- Apty · Hidden Costs of Poor ServiceNow Adoption · 2025

GCC-Native. Platform-Deep.

ACCOUNTABLE FOR RESULTS. AVERO IS MEA'S SERVICENOW AUTHORITY FOR ENTERPRISE TRANSFORMATION

Avero · May 2026

Most organisations have deployed ServiceNow. Few have moved beyond the system they received at go-live.

Avero · May 2026

HOW TO USE THIS PAPER

This paper does not reach a universal conclusion. The right governance architecture for a Microsoft-centric estate is different from the right answer for an organisation with ServiceNow as its system of record for core operations.

The Vertex Diagnostic exists precisely because honest advisory work requires making that determination first - before any platform is recommended.



5

SERVICE MATURITY STAGES ACROSS THE VERTEX FRAMEWORK. EVERY ENGAGEMENT GOVERNED. EVERY OUTCOME MEASURED

Avero · Vertex Framework™



<10%

OF ORGANISATIONS SCALING AI AGENTS HAVE GOVERNANCE IN PLACE AT SCALE

McKinsey 2026

SECTION 01

Executive Summary

GCC enterprises have deployed AI broadly. Most cannot demonstrate that the investment is governed, measurable, or compliant. The gap between deployment and defensibility is where credibility is lost, and where the most significant enterprise risk now lives.

\$58.3B

GCC digital transformation market 2025, growing to \$179.7B by 2030

90%

Organisations failing to realise full ServiceNow platform value

39%

MENA enterprises already running gen AI in production today

<10%

Scaling AI agents with governance in place - McKinsey 2026

GCC enterprises have deployed AI broadly. Most cannot demonstrate that the investment is governed, measurable, or compliant. The gap between deployment and defensibility is where credibility is lost, and where the most significant enterprise risk now lives. Governance that arrives after deployment is not governance. It is documentation of a risk that already materialised.

This paper synthesises three bodies of evidence: ServiceNow's Knowledge 2026 product announcements published in May 2026; McKinsey QuantumBlack's enterprise agentic platform architecture research published in April 2026; and the current state of GCC AI regulation across PDPL, QCB, and the Dubai AI Seal as each moves from aspiration to enforceable audit requirement. The central finding is that AI governance must live at the workflow execution layer. Both ServiceNow and McKinsey arrived at this conclusion independently. ServiceNow's commercial answer is AI Control Tower. McKinsey's architectural answer is a composable, compostable platform with governance embedded from day one.

The build versus buy decision resolves around three variables: the organisation's existing platform estate and integration backbone, its AI maturity and architectural readiness, and the urgency introduced by regulatory pressure and current agent sprawl. Avero's Vertex Diagnostic routes organisations through each dimension and maps them to one of five governance architecture lanes, each carrying a specific technology recommendation and defined engagement model. The diagnostic determines which lane applies before any platform recommendation is made, and platform preference is the output rather than the input.

AVERO'S CORE POSITION

For most GCC enterprise organisations with existing ServiceNow investment, AI Control Tower is the right governance architecture choice in 2026. The free year removes the financial barrier. The recommendation holds on one condition: that it is adopted within a structured delivery framework, architected with open protocols for compostability, and with year 2+ economics modelled explicitly before commitment is made.

The Vertex Diagnostic determines which of five governance architecture lanes applies to each organisation before any platform recommendation is made. That determination is the starting point, and platform preference is the output rather than the input.

The gap between deployment and defensibility is where the most significant enterprise risk now lives.

Avero · May 2026

THE THREE DECISION VARIABLES

Platform estate

What the organisation already operates and has invested in building.

AI maturity

Production-readiness across four foundational governance capabilities.

Regulatory urgency

The time available before the first audit artefact is formally required.



12-18 months

TO BUILD AN EQUIVALENT GOVERNANCE LAYER FROM OPEN-SOURCE COMPONENTS

Enterprise AI implementation benchmarks
Timeline indicative; varies by scope and capability



1,847

DOWNSTREAM REQUESTS AFFECTED BY A SINGLE MALICIOUS AGENT INSTRUCTION · K26 LIVE DEMO

ServiceNow K26 · May 2026

SECTION 02

The Problem: Deployment Without Defensibility

The most common AI deployment failure in the GCC is not technical. It is governance arriving after the AI is already live, leaving the organisation exposed at the precise moment the regulator or Board makes the request.

McKinsey QuantumBlack's April 2026 research frames the underlying dynamic as the gen AI paradox: broad adoption, limited bottom-line impact. Horizontal AI applications such as copilots, chatbots, and productivity tools are easy to deploy and scale, but sit loosely on top of the processes they are meant to transform. They deliver limited measurable value because they do not change how work is executed. Vertical agentic applications, embedded into specific business workflows, can be genuinely transformative, but only when governance infrastructure exists to make them defensible, accountable, and improvable over time.

ServiceNow's K26 keynote made the operational risk concrete. A live demonstration showed a malicious instruction embedded in an agent spreading to 1,847 downstream requests across a two-hour window, manipulating data and altering decisions, before being detected and terminated by AI Control Tower's real-time kill switch. This was not theoretical. It was a reproduction of a pattern already occurring in enterprise estates where AI agents operate without formal ownership, risk assessment, or audit trail.

The GCC enterprise context amplifies the global governance gap in three specific ways. First, the pace of AI deployment across sovereign institutions and large enterprise organisations in the UAE, Saudi Arabia, and Qatar has outrun the governance infrastructure in most organisations. Second, GCC-specific regulation is moving faster than most organisations appreciate, with PDPL enforcement live, QCB guidance moving from advisory to audit, and the Dubai AI Seal transitioning to a verifiable compliance requirement. Third, the master data foundation across many GCC enterprise estates remains fragmented, and AI does not fail because the model is wrong but because the data beneath it is.

MCKINSEY ON THE THREE ROOT CAUSES

Most agentic AI initiatives stall for one of three reasons: governance applied as a bolt-on rather than an architectural design principle; platforms not composable enough to evolve as the market matures; and production requirements treated as a later phase rather than an upfront constraint. All three have the same remedy: governance as a first-class concern from day one.

QuantumBlack, AI by McKinsey · April 2026

39%

OF MENA ENTERPRISES ARE ALREADY RUNNING GEN AI IN PRODUCTION

GSMA MENA Report 2025

THE REGULATORY TIMELINE

PDPL

Saudi Arabia. Enforcement live. Named AI asset ownership required.

QCB


Qatar. Advisory moving to audit. Financial services AI under review.

Dubai AI Seal

UAE. Aspirational standard moving to verifiable requirement.

EU AI Act


Cross-jurisdictional. Applies to any GCC org with European exposure.



\$500-\$1.2M

UPFRONT COST TO BUILD A CUSTOM ENTERPRISE AI GOVERNANCE LAYER – PLUS \$300-\$600K ANNUAL MAINTENANCE

*Credo AI Build vs Buy Analysis, 2026
Figures indicative; costs vary by scope and capability*



4

LIVE REGULATORY MANDATES AFFECTING GCC AI PROGRAMMES SIMULTANEOUSLY IN 2026

Avero · May 2026

SECTION 03

GCC Regulatory Considerations: PDPL, QCB, Dubai AI Seal, and EU AI Act

Regulatory pressure is no longer a future risk for GCC enterprise AI programmes. The mandates are live, the timelines are compressing, and the evidence requirements are moving from policy documents to audit artefacts.

PDPL: Saudi Arabia Personal Data Protection Law

PDPL enforcement is live and expanding in scope. Saudi organisations deploying AI in any process that touches personal data face documented obligations: named ownership of AI assets, completed risk assessments, data processing agreements, and subject rights management. AI Control Tower's Govern pillar, aligned to PDPL requirements out of the box, significantly accelerates compliance for ServiceNow-heavy estates. For other estates, the compliance framework must be built, a materially more costly and time-consuming path.

QCB: Qatar Central Bank AI Governance

QCB guidance on AI in financial services is moving from advisory to audit. Qatari financial institutions deploying AI agents in credit assessment, fraud detection, customer service, or operational processes face obligations mirroring the EU AI Act's high-risk classification requirements: transparency, human oversight, risk documentation, and continuous monitoring. The organisations most exposed are those that deployed AI during the exploratory phase without establishing governance infrastructure to evidence it.

Dubai AI Seal

The Dubai AI Seal is transitioning from an aspirational quality standard to a verifiable compliance requirement. Since October 2025, the Dubai Department of Finance has directed government entities to work exclusively with Seal-certified AI suppliers — making certification a procurement prerequisite for Dubai government AI projects. Organisations with governed, audit-ready AI infrastructure are materially better positioned to achieve and maintain certification. For ServiceNow-heavy estates, AI Control Tower provides that infrastructure out of the box.

EU AI Act: Applicable to GCC Organisations

EU AI Act enforcement is live and applicable to any GCC organisation with European operations, European customers, or data subjects located in the EU. The Act classifies AI systems by risk level and imposes obligations on both providers and deployers of high-risk AI. For GCC organisations with dual exposure, AI Control Tower's out-of-the-box frameworks cover both bodies of regulation simultaneously, making it the most efficient compliance path for organisations with cross-jurisdictional exposure.

THE REGULATORY IMPLICATION FOR THE BUILD VS. BUY DECISION

An organisation with an imminent PDPL audit or QCB review cannot afford the 12 to 18 months that building a governance layer from open-source components requires. For organisations facing imminent regulatory exposure, the build versus buy decision is effectively already made.

Live Now

ALL FOUR GCC REGULATORY MANDATES ARE IN FORCE. ASPIRATIONAL STANDARDS ARE NOW AUDIT REQUIREMENTS

Avero · May 2026

REGULATORY SNAPSHOT

PDPL

Saudi Arabia. Enforcement live. Named AI asset ownership required.

QCB

Qatar. Advisory to audit. Financial services AI under review.

Dubai AI Seal

UAE. Procurement requirement from Oct 2025. Certification requires demonstrable AI governance infrastructure.

EU AI Act

Cross-jurisdictional. Applies to any GCC org with European exposure.



6-10

WEEKS TO PDPL-ALIGNED, QCB-AWARE, DUBAI AI SEAL-MAPPED GOVERNANCE ON A SERVICENOW-HEAVY ESTATE



78%

OF BUSINESS EXECUTIVES LACK CONFIDENCE THEY COULD PASS AN INDEPENDENT AI GOVERNANCE AUDIT WITHIN 90 DAYS

Grant Thornton AI Impact Survey · 2026

SECTION 04

What Defensible AI Looks Like in Practice

The organisations successfully scaling agentic AI share one characteristic: they treated governance as a day-one architectural constraint, not a retrospective compliance exercise.

The organisations that have successfully governed AI deployment at scale share a common pattern: they treated governance as a day-one architectural constraint, not a retrospective compliance exercise. Before the first agent goes into production, ownership is documented, risk is assessed, an audit trail is established, and the human escalation path is defined for decisions that fall outside the agent's operating parameters. These are not post-deployment additions but the preconditions that determine whether the AI programme is defensible when the regulator or the Board makes the request.

The barrier is not technical complexity but the assumption that governance is something added later, once the use case is proven. That assumption is the most expensive mistake in AI programme management, because every agent deployed without governance infrastructure is a liability that compounds with every decision it makes, and the compounding is invisible until it materialises at exactly the wrong moment.

Seventy-eight percent of business executives lack confidence their organisation could pass an independent AI governance audit within 90 days, and AI deployment metrics are nine times less likely to have a named owner than standard business KPIs. These are not organisations that chose to defer governance but organisations that did not know governance needed to be present from the outset.

Defensible AI requires clean foundations. The master data across many GCC enterprise estates remains fragmented, and this is where AI programmes fail in practice. The model is rarely the problem. The data beneath it is. Every agent deployed into a fragmented data environment scales the wrong answer at a speed no human oversight can correct, and the compounding effect remains invisible until it materialises at precisely the moment the organisation can least afford it.

WHAT THIS LOOKS LIKE BEFORE GO-LIVE

The organisation that can answer four questions before an agent goes live is the organisation that can answer the Board twelve months later. Who owns this agent? What has it been assessed to do and not do? Where does the audit trail live? And who makes the call when it operates outside its parameters? Most organisations deploying AI in 2026 cannot answer all four. That is the governance gap in practice.

Avero · May 2026

9X

AI DEPLOYMENT METRICS ARE NINE TIMES LESS LIKELY TO HAVE A NAMED OWNER THAN STANDARD BUSINESS KPIs

ClearPoint / McKinsey State of AI Trust · 2026

FOUR PRE-CONDITIONS FOR DEFENSIBLE AI

Own it

Named ownership of every AI agent before it goes into production. No owner, no deployment.

Assess it

Completed risk assessment documenting what the agent can do, what it can access, and what happens when it fails.

Trace it

An audit trail from day one. Every decision the agent makes must be traceable before a regulator asks for it.

Escalate it

A defined human escalation path for every decision that falls outside the agent's operating parameters.

The gap

Fewer than 10% of organisations scaling AI agents have all four in place. That is the governance gap.



30%

ONLY 30% OF ORGANISATIONS HAVE GOVERNANCE READINESS IN PLACE DESPITE 58% HAVING AI IN ENTERPRISE-WIDE STRATEGY.

Deloitte State of AI · 2026



30+

ENTERPRISE INTEGRATIONS INCLUDING AWS, AZURE, GCP, SAP, ORACLE, WORKDAY AND MICROSOFT AGENT 365

ServiceNow K26 · May 2026

SECTION 05

What ServiceNow AI Control Tower is in 2026

AI Control Tower was introduced at K25 as a visibility dashboard. At K26 it was fundamentally reimagined into an active, cross-vendor control plane, bundled into every ServiceNow package by default and extended across 30 enterprise systems beyond ServiceNow's own ecosystem.

The architectural significance of this evolution is often underestimated. The original AI Control Tower was an observability tool: it showed what AI assets existed and how they were performing. The K26 version is an active governance infrastructure. It does not merely observe agents. It governs them in real time, constrains them within defined permissions, and can terminate them when they operate outside their intended scope. Most competing governance tools operate at the observation layer. AI Control Tower operates at the execution layer.

Four announcements at K26 materially expanded AI Control Tower's addressable scope. The Microsoft Agent 365 integration extends governance to Copilot Studio and Azure AI Foundry agents, covering the Microsoft infrastructure most GCC enterprises already operate and positioning ServiceNow as governing Microsoft's own agents at the execution layer that Microsoft cannot reach natively. The AI Gateway, built on the open MCP protocol, governs all agent-to-tool transactions in real time and preserves multi-vendor flexibility. Project Arc with NVIDIA extends governance to autonomous desktop agents, and the Amazon Bedrock AgentCore integration allows organisations to build agents on AWS and govern them through AI Control Tower as the unified control plane.

AI Control Tower is included in every ServiceNow product and package, no longer an add-on SKU. It is available free for one year, a programme valued at two million dollars. This is a deliberate land-and-expand strategy: remove the financial barrier, establish platform dependency, then capture the renewal. Enterprise leaders should model year two-plus economics explicitly before committing. The free year is a genuine opportunity but it is not a permanent position.

THE ARCHITECTURAL CLAIM THAT MATTERS

ServiceNow's governance is not model-level or data-level. It is execution-level. An AI agent that can initiate an approval workflow, trigger a playbook, or update a catalogue record, with a full audit trail, is a fundamentally different governance surface than one that can only query a database. Twenty-two years of enterprise operational data gives ServiceNow the business context to govern what agents do, not just what they know.

ServiceNow Knowledge 2026 · May 2026

Most governance tools observe agents. AI Control Tower governs them.

Avero · May 2026

FIVE FUNCTIONAL PILLARS

Discover

30 new enterprise integrations covering AWS, Azure, GCP, SAP, Oracle, Workday, and Microsoft Agent 365.

Observe

Live runtime monitoring of agent behaviour via the Traceloop acquisition. Real-time visibility into what agents are doing.

Govern

Five risk frameworks aligned to NIST AI RMF, EU AI Act, PDPL, QCB, and Dubai AI Seal standards, out of the box.

Secure

Identity access governance via Veza integration. 30 billion permissions mapped across every human and non-human identity.

Measure

Cost tracking and ROI dashboards. \$462.8K AI value tracked with a 33% improvement in 30 days across 1,884 assets governed.



22 years of enterprise operational data. The context to govern what agents do, not just what they know.

ServiceNow K26 · May 2026



3

LIVE FINANCIAL INSTITUTION DEPLOYMENTS UNDERPINNING THE MCKINSEY FRAMEWORK

McKinsey QuantumBlack · April 2026

SECTION 06

The McKinsey Architecture Framework: Composable and Compostable

McKinsey QuantumBlack's April 2026 framework provides a vendor-agnostic architecture for enterprise agentic platforms that is the most accessible thinking available in the market. It is the essential counterweight to any single-platform governance narrative, including AI Control Tower's.

The McKinsey framework is grounded in three live financial institution deployments: a European bank automating credit application workflows; a large financial services organisation building a digital factory of agents for software delivery; and a global bank reinventing its software development lifecycle with end-to-end automated agent workflows. The findings are directly applicable to GCC enterprise contexts, where financial services and sovereign institutions face comparable complexity.

Enterprise agentic platforms should be both composable, built from modular components that can be assembled into complex systems, and compostable, designed so individual components can be replaced without requiring a complete architectural redesign. This principle is the primary safeguard against the platform lock-in risk inherent in any governance architecture choice, including AI Control Tower.

In practice, McKinsey recommends a build versus buy versus partner framework applied component by component: buy where clear market solutions exist; partner where viable solutions are likely to emerge; build selectively only where genuine differentiation is possible. The majority of components should be purchased. Building is reserved for configuration, integration, and orchestration.

Open protocols are the primary mechanism for preserving multi-vendor flexibility. Two protocols are emergent and significant: Agent2Agent enables direct communication between agents across platforms; Model Context Protocol enables secure, governed access by agents to external tools and data at runtime. ServiceNow's AI Gateway at K26 is built on MCP, a positive signal that even a platform-centric architecture can embrace open standards.

MCKINSEY ON BUILD VS. BUY

In practice, most agentic platform components are purchased, with building reserved primarily for configuration, integration, and orchestration. When partnering or building a solution likely to become available in the market in the medium term, CIOs should ensure that the short-term benefit outweighs the temporary technical debt it incurs.

QuantumBlack, AI by McKinsey · April 2026

McKinsey QuantumBlack

APRIL 2026 · ENTERPRISE AGENTIC PLATFORM ARCHITECTURE FRAMEWORK

A vendor-agnostic architecture framework for enterprise agentic platforms.

FOUR PRODUCTION-READINESS CAPABILITIES

- Agentic evaluation**
Testing agent behaviour before and continuously after deployment.
- Agent discoverability**
Registries of what agents exist and what they are authorised to do.
- Memory management**
Context continuity across sessions, reducing cost and improving governance quality.
- Feedback loops**
Continuous improvement mechanisms that keep governance current against an evolving operational environment.



Day One

PRODUCTION READINESS MUST BE AN ARCHITECTURAL CONSTRAINT, NOT A LATER PHASE



5

GOVERNANCE ARCHITECTURE OPTIONS
EVALUATED FOR GCC ENTERPRISE ESTATES

Avero Market Assessment · May 2026

SECTION 07

The Competitive Landscape: Choosing the Right Governance Architecture

The question is not which governance platform is best in the abstract. It is which governance architecture is most appropriate for a specific organisational estate. Five profiles shape that answer: existing ServiceNow footprint, data and integration maturity, hyperscaler relationships, tolerance for vendor dependency, and internal capacity to build and sustain platform engineering.

Microsoft Purview and Agent 365

Microsoft Purview and Agent 365 are already present in most GCC enterprises through M365 investment. Purview governs data: classification, lineage, and compliance policy enforcement. The architectural gap is workflow execution governance. Microsoft's stack does not natively understand the approval chains, escalation logic, and business rules that determine whether an agent action is appropriate in a given business context. ServiceNow's K26 extension of AI Control Tower into Microsoft Agent 365 directly addresses this gap.

Boomi, MuleSoft, and Integration-Platform-Led Governance

Boomi and MuleSoft are not competitors to AI Control Tower. They are prerequisites for it in any estate where master data governance is the primary challenge. An AI agent trained on fragmented or inconsistent data scales the wrong answer at a speed no human oversight can correct. For integration-platform-led estates, the governance conversation must begin with the data foundation. AI Control Tower activates once the foundation is solid.

Hyperscalers: AWS Bedrock and Google Vertex

AWS Bedrock AgentCore and Google Vertex AI provide strong model governance and agent runtime capabilities at the infrastructure layer. The gap is business-process governance: they do not understand the approval workflows, compliance obligations, and operational context that determine whether an agent action is appropriate in a GCC enterprise environment. The recommended path for hyperscaler-heavy estates is complementarity rather than replacement.

Purpose-Built and Open-Source Governance

Purpose-built vendors including Aporia, Fiddler AI, Arthur AI and Phoenix by Arize AI offer modular, vendor-neutral governance at significantly lower lock-in risk. Their limitation is scope: they operate at the model and output layer and do not govern workflow execution. This approach requires dedicated platform engineering investment that most GCC enterprises do not have internally, but it is the most architecturally defensible position on vendor dependency for organisations with the capacity to sustain it.

ESTATE-FIRST PRINCIPLE

Leading with platform preference before understanding the estate is the most common advisory failure Avero encounters. A Boomi-centric organisation with no ServiceNow footprint needs a fundamentally different answer than an enterprise with ServiceNow as its system of record for core operations. That determination must come first.

Estate First

ALWAYS. BEFORE ANY PLATFORM
RECOMMENDATION IS MADE

Avero Advisory Position · 2026

KEY COMPARISON DIMENSIONS

Execution-layer governance

SN AICT only - native, including identity governance across human and non-human identities. Others: data or model layer only.

GCC regulatory frameworks

SN AICT out of box. MS EU/US focus. Others require build.

Cross-vendor agent discovery

SN AICT 30+ integrations. MS centric. OSS requires build.

Real-time kill switch

SN AICT K26 proven. MS policy-based only. OSS custom.

Vendor lock-in risk

SN AICT and MS: High. Boomi: Medium. Purpose-built: Low. OSS: Lowest.



Year Two

THAT IS WHEN THE PLATFORM DEPENDENCY
BECOMES A COMMERCIAL OBLIGATION



*Urgency first.
Estate second.
Maturity third.*

Avero · Vertex Diagnostic Protocol

SECTION 08

The Vertex Diagnostic and Five Outcome Lanes

Before recommending any governance platform, Avero runs every organisation through the Vertex Diagnostic: a structured three-stage assessment that establishes urgency, estate reality, and AI maturity before any architecture recommendation is made.

Stage 1: Governance Urgency

Three questions are sufficient: Are AI agents running in production today without centralised governance? Does the organisation have visibility of its non-human identities and what they can access? What is the regulatory deadline pressure across PDPL, QCB, Dubai AI Seal, and EU AI Act? High urgency compresses the entire diagnostic and accelerates the architecture decision.

Stage 2: Estate and Integration Reality

The most consequential routing decision in the diagnostic. Leading with AI maturity questions before understanding the estate is the most common advisory failure. Five questions establish the lane: the dominant workflow platform, the integration backbone and master data maturity, the degree of AI vendor fragmentation, the primary cloud infrastructure, and the organisation's tolerance for platform vendor dependency.

Stage 3: AI Maturity and Architectural Readiness

The final stage refines the lane and determines the investment level and Avero engagement model. McKinsey's four production-readiness capabilities provide the assessment framework: agentic evaluation, agent discoverability, memory management, and feedback loops. Where these foundations are weak, the Avero engagement begins there before any platform is activated.

Platform preference is the output. Not the input.


Avero · Vertex Framework™

DIAGNOSTIC PRINCIPLES

- Urgent first**
Governance emergencies compress the diagnostic. Regulators do not wait.
- Estate second**
Understand what the organisation already runs before evaluating platforms.
- Maturity third**
AI maturity shapes the engagement model, not the lane assignment.

FIVE OUTCOME LANES · VERTEX FRAMEWORK™ ARCHITECTURE DECISION

Lane	Client Profile	Verdict	Governance Architecture	Avero Engagement
1 · SN-heavy	Strong ServiceNow footprint, existing workflow investment	Buy - AICT as control plane	AICT as enterprise control plane	Accelerate & govern via Vertex Framework
2 · MS-heavy	M365, Azure AI Foundry, Copilot Studio	Hybrid - Purview + AICT execution layer	AICT governs MS agents at execution layer	Architect the Purview / AICT boundary
3 · Integration-led	Boomi, MuleSoft, strong MDM	Data-first, selective AICT	AICT activates where SN workflows exist	Data Integrity leads; AICT scoped after
4 · Hyperscaler	AWS Bedrock, Google Vertex, multi-cloud	Bedrock/Vertex + AICT overlay	AICT at the business-process governance layer	Complementarity architecture design
5 · Greenfield	No incumbent platform, build-first preference	Build - OSS stack	AICT not recommended at this stage	Platform architect & OSS delivery partner



AI CONTROL TOWER

AICT CAN BE THE RIGHT ANSWER IN FOUR OF FIVE LANES. IN LANE 5 IT IS NOT. THE VERTEX DIAGNOSTIC EXISTS TO DETERMINE WHICH LANE APPLIES BEFORE ANY PLATFORM CONVERSATION BEGINS



3

NON-NEGOTIABLE ADVISORY PRINCIPLES - THE BASIS OF EVERY AVERO ENGAGEMENT

Avero Advisory Position · 2026

SECTION 09

Avero's Point of View and Recommended Actions

Three principles apply regardless of which governance architecture lane an organisation occupies. They are Avero's non-negotiable advisory positions and the basis on which every Avero engagement is structured.

Principle One Govern at the execution layer. Not after the fact.


Governance applied after deployment is documentation of a risk that already materialised. Both McKinsey's research and ServiceNow's K26 position reach the same conclusion: governance must be embedded in the workflow execution layer from the first agent in production. Ownership documentation, risk assessments, and compliance evidence must exist before the AI goes live. The organisation that waits until regulators demand them is already exposed.

Principle Two Composable and compostable. Design for replacement.

Every governance architecture choice carries a lock-in cost. Acknowledging that cost and designing around it is the difference between a platform commitment made with full awareness and a dependency discovered at renewal. Open protocols at the integration layer preserve optionality even within a ServiceNow-centric architecture. Year 2+ economics must be modelled before the free year begins, not at the end of it.

Principle Three Production readiness from day one. Not as a retrofit.

McKinsey's most important finding: the organisations that successfully scale agentic AI treat production requirements as architectural constraints from the outset. Agentic evaluation, agent discoverability, memory management, and feedback loops must be established from the first agent in production. Activating a governance platform on inadequate foundations produces governance theatre, not governance.



<1/3

OF ORGANISATIONS CAN TIE THE VALUE OF AI TO THEIR ORGANISATION'S FINANCIAL GROWTH

Forrester · 2026 Technology & Security Predictions

AVERO'S GOVERNANCE ARCHITECTURE VERDICT

Activate AI Control Tower. Architect for compostability. Govern what comes after.

For most GCC enterprise organisations with existing ServiceNow investment, AI Control Tower is the right governance architecture choice in 2026. The free year removes the financial barrier. The recommendation holds on one condition: adopted within the Vertex Framework, architected with open protocols, and with year 2+ economics modelled explicitly before commitment.

THREE ACTIONS EVERY ORGANISATION SHOULD TAKE NOW

01

Commission an AI Discovery Session

Establish where AI is running, who owns it, and where the governance gaps are. The output is a structured evidence base that gives leadership a clear picture of exposure before any platform decision is made.

02

Run the Vertex Diagnostic before selecting a platform

Urgency first. Estate second. Maturity third. The most common advisory failure in this market is recommending a platform before understanding what the organisation actually has. The diagnostic takes days, not weeks. The cost of skipping it is measured in years.

03

Model year 2+ economics before year 1 begins

AI Control Tower's free year is a commercial strategy. The organisations that benefit are those who enter with a clear view of the commitment they are making and the value they expect to demonstrate before the first renewal conversation.

Your AI is live. Is it defensible?

Avero's AI Discovery Session establishes where your AI is running, who owns it, and where the governance gaps are. It is a diagnostic, not a sales call. The output is yours, regardless of what follows.

Book an AI Discovery Session

[AVERO.WORKS/CONTACT](https://avero.works/contact)

Vertex Framework is a trademark of Avero · All figures indicative. Avero makes no representation as the completeness of third-party research cited herein