WHITE PAPER · APRIL 2026

Agentic Workforce Management™

A Framework for Acquiring, Onboarding, Managing, and Offboarding AI Agents in the Enterprise

AUTHOR

Michael Beygelman

TITLE

Co-Founder & CEO, Insygna Corporation

PUBLISHED

April 2026

ABSTRACT

Autonomous AI agents are now a meaningful component of the enterprise workforce. They take actions, make decisions, and trigger workflows without human involvement in each transaction. No established management discipline exists for governing this new workforce tier. This paper introduces Agentic Workforce Management™ (AWM™), a formal framework for acquiring, onboarding, managing, and offboarding AI agents, and argues that organizations must build this competency now, before scale makes ungoverned agent deployments a material operational and regulatory liability.

Insygna

Table of Contents

Executive Summary

The enterprise workforce is changing in ways that existing management disciplines were not designed to handle. Autonomous AI agents are being deployed inside organizations at scale, taking actions, making decisions, triggering workflows, and communicating with employees and customers without direct human involvement in each transaction. By the end of 2025, Gartner estimated that more than 15% of day-to-day business decisions would be made autonomously by AI systems. By 2028, that figure is projected to reach 33%.

This transformation creates a governance gap that is structural, not procedural. Organizations have spent decades building frameworks for managing human workforces, and later for managing extended and contingent workforces. None of those frameworks were designed for a workforce participant that operates at machine speed, requires no salary, spawns sub-agents dynamically, and leaves no inherent identity trace unless one is deliberately built into the infrastructure.

This paper introduces **Agentic Workforce Management™ (AWM™)** as a formal discipline: the set of organizational practices, policies, and technical infrastructure required to safely acquire, onboard, manage, and offboard autonomous AI agents operating in enterprise environments. AWM™ borrows from established workforce management practice, extends it with requirements specific to agent behavior, and provides leaders with a framework they can begin to implement now.

The urgency is not speculative. The EU AI Act, with its enforcement deadline of August 2, 2026, places binding obligations on organizations deploying AI agents in high-risk categories. Non-compliance carries fines of up to EUR 35 million or 7% of global annual turnover, whichever is higher. But regulatory compliance is the floor, not the ceiling. The organizations that will lead in the agentic era are those that develop AWM™ as an operational competency.

This paper is written for senior leaders responsible for how work gets done: Chief Human Resources Officers, Chief People Officers, Chief Information Officers, heads of extended workforce programs, and the operations and procurement leaders who manage the systems through which contingent work flows. It assumes familiarity with enterprise workforce management concepts and does not require a technical background in AI systems.

> **AWM™ DEFINED**
>
> *Agentic Workforce Management™ is the organizational discipline governing the full lifecycle of autonomous AI agents operating within enterprise environments: from sourcing and credentialing through deployment and oversight to decommissioning and audit retention.*

SECTION 01

The Composition of the Modern Enterprise Workforce

Three Tiers of Work

For most of the twentieth century, the enterprise workforce had two meaningful categories: employees and contractors. The distinction shaped everything from HR policy to payroll systems to compliance obligations. Beginning in the 1990s, the rise of the contingent and extended workforce added complexity. Vendor Management Systems (VMS) emerged as a category specifically to govern the procurement, onboarding, and management of non-employee workers at scale. Organizations built infrastructure to bring discipline to a workforce tier that traditional HRIS systems were never designed to handle.

That evolution took decades. The transition to a three-tier workforce, adding autonomous AI agents as a genuine participant class alongside employees and contingent workers, is happening in years.

As of 2026, AI agents are operating inside enterprise environments performing work that previously required human judgment: screening resumes, scheduling interviews, responding to employee inquiries, processing invoices, drafting communications, analyzing data, and triggering downstream system actions. In a meaningful operational sense, they are doing work. The question is whether the organization is managing them accordingly.

The Scale of Agent Deployment

> **By 2028, 33% of enterprise software applications will include agentic AI capabilities, up from less than 1% in 2024.**
>
> *Source: Gartner, Inc., "Predicts 2025: Agentic AI," October 2024. gartner.com/en/newsroom/press-releases/2024-10-21-gartner-predicts-agentic-ai*

> **82% of organizations plan to integrate AI agents within the next three years, with 24% already actively deploying them.**
>
> *Source: Salesforce, "State of IT: Agentic AI," 2025. salesforce.com/news/stories/state-of-it-ai-agents/*

These projections reflect procurement decisions already underway. Enterprise software vendors including Salesforce, ServiceNow, Microsoft, SAP, and Workday have all announced agentic capabilities embedded in their core platforms. Organizations that have not yet encountered an AI agent acting on their behalf are the exception, not the rule.

The Analogy to Contingent Workforce Management

The contingent workforce analogy is instructive, though it has limits. When organizations began deploying contract workers at scale in the 1990s, the initial instinct was to treat them as a procurement matter, not a workforce matter. They were handled through accounts payable, not HR. The absence of governance frameworks created compliance exposure: misclassification liability, inconsistent access controls, audit gaps, and no systematic offboarding when engagements ended.

The VMS category emerged because organizations eventually recognized that contingent workers, despite not being employees, were doing work that carried organizational risk. They needed to be sourced through defined channels, vetted before onboarding, granted appropriate access, managed during engagement, and properly offboarded when the engagement concluded.

> **The global contingent workforce management market was valued at USD 215.7 billion in 2023 and is expected to exceed USD 450 billion by 2031.**
>
> *Source: Zion Market Research, "Contingent Workforce Management Market," 2024. zionmarketresearch.com/report/contingent-workforce-management-market*

AI agents present an analogous governance challenge, with higher stakes. A contingent worker who acts outside their defined scope creates liability. An AI agent that acts outside its defined scope can do so at machine speed, across multiple systems simultaneously, with no human judgment applied to each individual action. The risk surface is orders of magnitude larger. The management discipline required is correspondingly more rigorous.

SECTION 02

Why Existing Governance Frameworks Fail

Frameworks Built for a Different Operational Reality

Enterprise governance frameworks for AI, where they exist, were designed for a specific technical architecture: a human submits a prompt, a model returns a response, a human reviews and acts on that response. This is the architecture of AI as a tool. It is not the architecture of AI as a workforce participant.

Agentic AI operates differently. An agent receives a goal, decomposes it into tasks, selects and uses tools to accomplish those tasks, evaluates its own outputs, and proceeds to the next step, all without a human reviewing each intermediate action. Multi-agent systems add further complexity: an orchestrating agent assigns tasks to specialized sub-agents, each of which may spawn additional agents or call external services. The action chain that results may involve dozens of discrete decisions and system interactions, none of which were individually authorized by a human.

Current enterprise governance frameworks have four structural gaps when applied to this architecture.

The Four Structural Gaps

GAP 1: NO MULTI-AGENT ORCHESTRATION GOVERNANCE

Enterprise AI policies typically govern the deployment of specific AI systems: a named tool, a defined use case, an approved vendor. They are not designed to govern the dynamic, runtime composition of agent networks where a human approves a top-level agent but that agent then assembles and directs subordinate agents whose individual behaviors were never separately reviewed or approved.

This creates accountability gaps that standard policy frameworks cannot close. If a sub-agent takes an action that causes harm, existing governance structures provide no mechanism to trace that action back to an authorization decision. The organization cannot answer the most basic compliance questions: who approved this agent to take this action, under what authority, and with what constraints.

GAP 2: NO ACCOUNTABILITY MECHANISMS FOR AUTONOMOUS ACTIONS

Human workforce accountability is built on identity. When an employee takes an action, they are authenticated into the systems they access, their actions are logged against their identity, and that record can be retrieved for audit or investigation. The accountability chain works because the identity layer is robust.

Most enterprise AI deployments today share API keys, service accounts, or access tokens across multiple agents and agent instances. An action taken by one agent instance is indistinguishable in audit logs from an action taken by any other agent sharing the same credentials. When something goes wrong, organizations cannot attribute the action to a specific agent instance, verify that the agent was operating within its defined scope, or demonstrate to a regulator that appropriate oversight was applied.

> **Only 24% of organizations report having visibility into what actions their AI systems are taking across enterprise systems.**
>
> *Source: IBM Institute for Business Value, "CEO Study: The Urgency of AI Governance," 2025. ibm.com/thought-leadership/institute-business-value/en-us/c-suite-study/ceo*

GAP 3: THE LATENCY PROBLEM IN HUMAN OVERSIGHT

Regulatory frameworks and internal governance policies frequently require human oversight of AI decisions. The EU AI Act mandates human oversight for high-risk AI systems (Article 14). Many organizations have implemented approval workflows for AI outputs before those outputs are acted upon.

This approach works when AI operates at human speed. It does not work when AI agents are processing thousands of transactions per hour. Human review becomes a bottleneck that either eliminates the operational benefit of the agent or, more commonly, becomes a rubber-stamp exercise where reviewers approve items without genuine evaluation because volume makes actual review impossible.

Effective oversight of agentic AI cannot be designed as a per-action human review process. It must be designed as an architecture: behavioral boundaries, trust scopes, anomaly detection, and escalation triggers that apply at the system level.

GAP 4: NO AGENT IDENTITY AND CREDENTIALING STANDARD

Human workforce management depends on a foundational assumption: every participant has a verified identity. Employees are onboarded through HR processes that establish their identity, their role, and their authorized scope. Contingent workers go through similar processes, managed through VMS platforms that maintain their credentials, engagement terms, and access rights.

No equivalent standard exists for AI agents. Agents are deployed without persistent, verifiable identities that travel with them across systems. They carry no credentials that specify their authorized scope in a machine-readable format that enterprise systems can enforce at runtime. They cannot be queried for their behavioral history, their current trust status, or their assigned permissions the way a human worker can be queried through an HRIS.

> **Machine identities, service accounts, API keys, tokens, and certificates, already outnumber human identities by 45 to 1 in the average enterprise. AI agents will accelerate that ratio substantially, without any corresponding increase in identity governance infrastructure.**
>
> *Source: CyberArk, "Identity Security Threat Landscape Report," 2024. cyberark.com/resources/threat-research-blog/2024-identity-security-threat-landscape-report*

This is not a configuration problem. It is an infrastructure gap. The identity layer for the AI workforce does not exist in most enterprise environments. Building it is the foundational requirement for any serious AWM™ program.

> **THE CORE INSIGHT**
>
> *Identity and Access Management (IAM) became non-negotiable infrastructure for the human workforce because organizations recognized that you cannot govern access, accountability, or compliance without persistent, verifiable identity. The AI workforce requires the same foundational layer.*

SECTION 03

The AWM™ Lifecycle

Agentic Workforce Management™ organizes the governance of AI agents around four sequential lifecycle stages, each with defined organizational practices, technical requirements, and accountability mechanisms. The model is deliberately parallel to mature contingent workforce management practice, because the governance logic is similar even when the technical implementation differs substantially.

The diagram below illustrates the AWM™ lifecycle and the identity layer at its center. Each stage depends on the identity infrastructure: without persistent, verifiable agent identity, acquisition cannot be verified, onboarding cannot be completed, management cannot be attributed, and offboarding cannot be certified.

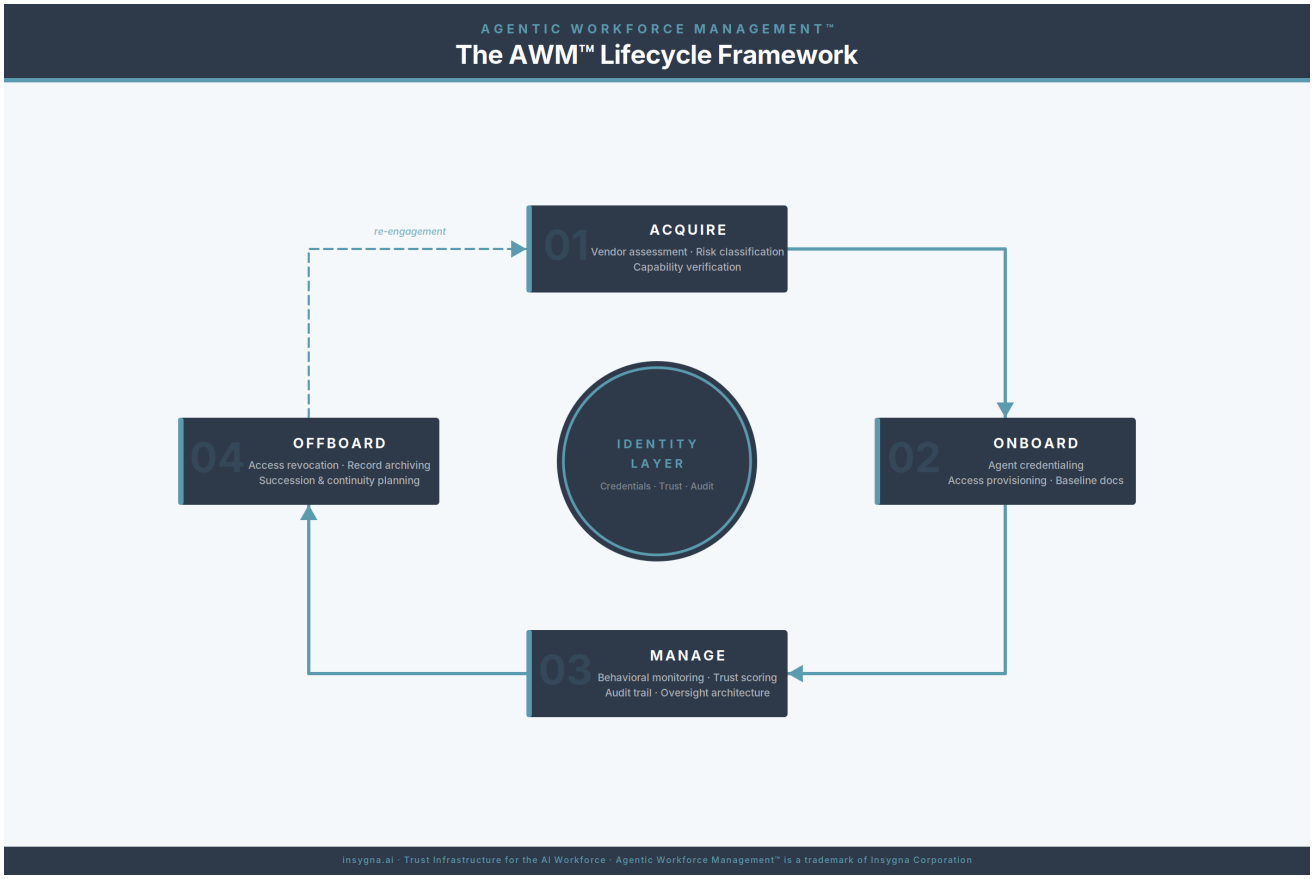

Figure 1. The AWM™ Lifecycle. The identity layer at the center anchors all four stages. The dashed arc represents re-engagement, where an offboarded agent may be re-credentialed for a new engagement under the same governance structure.

3.1 Acquisition

Acquisition is the process by which an organization evaluates, sources, and approves an AI agent before that agent enters the enterprise environment. In contingent workforce management, acquisition includes vendor qualification, worker screening, contract negotiation, and engagement setup. In AWM™, the equivalent processes address a different set of risk factors.

AGENT SOURCING AND VENDOR ASSESSMENT

Before deploying an agent, the organization must establish where it comes from, who built it, what it was trained on, and under what terms it operates. A structured agent vendor assessment should address: the vendor's AI development practices and model documentation; data handling and residency terms; explainability and audit capabilities; the vendor's own compliance posture under applicable regulations; and contractual accountability terms for agent behavior.

The EU AI Act requires providers of high-risk AI systems to maintain technical documentation sufficient to demonstrate compliance. Organizations deploying third-party agents bear responsibility for ensuring that documentation exists and is accessible. (EU AI Act, Article 11.)

CAPABILITY VERIFICATION AND SCOPE DEFINITION

Before deployment, the organization must define, in explicit terms, what the agent is authorized to do, the AWM™ equivalent of an employment contract or statement of work. The scope definition specifies: the tasks the agent is authorized to perform; the systems it is authorized to access; the data it is authorized to read and write; the external services it is authorized to call; and the circumstances under which it must escalate to a human.

Scope definition that exists only in a document is insufficient. The defined scope must be encoded in technical controls enforced at the infrastructure level, not merely stated in policy. An agent that has been told it should not access compensation data must be technically prevented from doing so.

RISK CLASSIFICATION

Not all agents carry the same risk profile. AWM™ programs should implement a risk classification framework that determines the appropriate level of oversight, the approval pathway for deployment, and the ongoing monitoring requirements based on the agent's scope and decision-making authority. The EU AI Act provides a useful reference framework, classifying AI systems by risk level with corresponding compliance obligations. (EU AI Act, Articles 5–16.)

3.2 Onboarding

Onboarding in AWM™ is the process of establishing an agent's operational identity, provisioning its access, and documenting its entry into the organizational environment. This is where the technical infrastructure of AWM™ is most clearly visible, and where most organizations currently have no process at all.

AGENT CREDENTIALING

Every agent deployed in an enterprise environment should have a persistent, verifiable identity credential before it takes any action. This credential is not an API key shared across instances. It is a unique, auditable identifier that travels with the agent across the systems it accesses, enabling every action to be attributed to a specific credentialed agent instance, the AWM™ equivalent of employee onboarding into an HRIS.

ACCESS PROVISIONING

Access provisioning for agents should follow the principle of least privilege, granting only the permissions required for the agent to accomplish its defined tasks. In the context of agentic AI, least privilege is a governance mechanism as much as a security control. Tight access provisioning limits the blast radius of agent misbehavior in a way that no amount of monitoring can replicate.

BASELINE DOCUMENTATION

At the time of onboarding, the organization should document the agent's expected behavioral baseline: the types of actions it is expected to take, the frequency and volume of those actions, the data it typically reads and writes, and the escalation conditions under which it should defer to a human. This documentation serves two functions: it provides the reference point for ongoing monitoring, and it provides the evidentiary basis for compliance demonstration.

> **Organizations with documented AI behavioral baselines detect anomalous agent behavior 3.4 times faster than those relying on reactive monitoring alone.**
>
> *Source: Capgemini Research Institute, "Harnessing the Value of Generative AI," 2024. capgemini.com/insights/research-library/generative-ai/*

3.3 Management

The management phase encompasses ongoing oversight of agents during their operational deployment. This is the most complex and most underdeveloped component of AWM™ practice, and where the limitations of frameworks designed for human oversight are most acutely felt.

BEHAVIORAL MONITORING AND DRIFT DETECTION

Agents can behave differently over time. Model updates, shifts in the data they encounter, prompt injection attacks, and gradual scope expansion through accumulated precedent can all cause divergence from a documented baseline. Behavioral monitoring in AWM™ is a technical capability: automated comparison of agent actions against defined behavioral parameters, with anomaly flagging and escalation triggers built into the oversight architecture.

TRUST SCORING

An agent's trust status should not be binary. AWM™ programs should implement dynamic trust scoring that reflects the agent's behavioral record, the consistency of its actions with its defined scope, and the outcomes of any prior incidents. Trust scoring enables calibration of human oversight investment to actual risk: high-trust agents in stable, well-defined scopes require less intensive monitoring than newly onboarded agents or agents operating in high-risk classifications.

AUDIT TRAIL MAINTENANCE

Every action taken by an agent should be recorded in an immutable audit trail tied to the agent's identity credential, capturing the agent's identity, timestamp, systems accessed, data read and written, and the outcome. The EU AI Act requires providers and deployers of high-risk AI systems to maintain logs sufficient to enable post-hoc review and accountability attribution. (EU AI Act, Article 12.)

> **Enterprises that implement AI audit logging report 67% faster resolution time in AI-related incidents compared to those without formal logging infrastructure.**
>
> *Source: Accenture, "Responsible AI: From Principles to Practice," 2025. accenture.com/us-en/insights/technology/responsible-ai*

HUMAN OVERSIGHT ARCHITECTURE

Human oversight of agentic AI must be designed as an architecture, not a per-action review process. The organization must define: which categories of agent action require pre-authorization by a human; which categories are within autonomous scope but trigger notification; which categories are fully autonomous within defined boundaries; and what conditions trigger escalation from autonomous to supervised operation.

3.4 Offboarding

Offboarding is the most frequently neglected component of AWM™ practice, and the one where the consequences of negligence are most directly analogous to well-understood problems in contingent workforce management. Contingent workers who are not properly offboarded retain access to enterprise systems after their engagement has concluded. AI agents that are not properly offboarded present the same class of problem, with a larger potential impact.

ACCESS REVOCATION

When an agent's engagement concludes, its credentials should be revoked and its access terminated across all systems. This requires that the organization maintains a complete, current registry of all systems the agent had access to, all credentials issued, and all integrations through which the agent could act. Organizations that onboarded agents without maintaining this registry cannot revoke access they do not know was granted.

BEHAVIORAL RECORD ARCHIVING

At offboarding, the agent's full behavioral record, including its audit trail, trust score history, incident reports, and scope documentation, should be archived in a manner that preserves its accessibility for post-engagement review. Regulatory frameworks including the EU AI Act specify retention periods for AI system logs. Beyond regulatory compliance, organizations have a practical interest in retaining behavioral records to support investigations that may arise after an agent's engagement has concluded.

SUCCESSION AND CONTINUITY

Offboarding must be coordinated with either a replacement agent deployment or a transition of the process back to human management. If the agent was the primary mechanism for executing a business process, its offboarding without succession planning creates operational continuity risk that is distinct from, but equally important to, the governance risks addressed elsewhere in the AWM™ lifecycle.

> **THE OFFBOARDING IMPERATIVE**
>
> *A study of enterprise IT security incidents found that 67% of unauthorized access events involving former contingent workers were attributable to failure to revoke access at offboarding. The same failure mode applies directly to AI agents. (Source: CyberArk, "Identity Security Threat Landscape Report," 2024. cyberark.com/resources/threat-research-blog/2024-identity-security-threat-landscape-report)*

SECTION 04

Trust Infrastructure as the Foundation

Why AWM™ Requires an Identity Layer

The AWM™ lifecycle depends, at every stage, on a capability that most organizations do not currently have: the ability to assign, maintain, and enforce a persistent, verifiable identity for each AI agent operating in their environment. Acquisition requires knowing what agents exist and what they are authorized to do. Onboarding requires issuing credentials recognized by enterprise systems. Management requires attributing actions to specific credentialed agents. Offboarding requires revoking specific credentials and archiving behavioral records. None of these capabilities exist without an identity layer.

This is precisely the problem that IAM solved for the human workforce. Before robust IAM infrastructure, organizations could not reliably answer: who accessed this data, when, and why? After IAM became standard enterprise infrastructure, those questions became answerable. AWM™ requires the same capability for the AI workforce.

What Agent Identity Infrastructure Provides

An agent identity and credentialing infrastructure provides four core capabilities. **Persistent identity**, a unique, auditable identifier for each agent instance that does not change across system interactions, recognized by enterprise systems the same way an employee's user identity is recognized. **Scope encoding**, the agent's authorized scope embedded in the credential itself, making it machine-readable and enforceable at the system level at runtime. **Behavioral record**, a persistent record of the agent's action history tied to its identity, enabling accountability attribution and trust score calculation. **Lifecycle management**, the ability to issue, modify, suspend, and revoke agent credentials through a managed process with full audit records of each lifecycle event.

Agent identity infrastructure is not a replacement for existing enterprise systems. It is a layer that integrates with existing HRIS, VMS, IAM, and SIEM systems, providing the agent-specific identity and behavioral record capabilities those systems were not designed to provide. For organizations with mature VMS deployments, the AWM™ identity layer extends the same discipline that VMS platforms brought to contingent workforce management into the agent workforce tier.

SECTION 05

Regulatory and Compliance Context

The EU AI Act

The EU AI Act (Regulation (EU) 2024/1689) represents the most comprehensive binding regulatory framework for AI systems currently in force. AI systems used in employment, worker management, and access to self-employment, including recruiting, selection, promotion, and performance evaluation, are explicitly classified as high-risk under Annex III, Section 4. AI agents operating in HR and talent management contexts fall squarely into this classification.

For high-risk systems, the Act requires: implementation of risk management systems throughout the lifecycle; data governance practices for training and operational data; technical documentation sufficient for compliance verification; automatic logging of events; transparency measures for users; human oversight measures that enable human intervention; and accuracy, robustness, and cybersecurity requirements. (EU AI Act, Articles 9–15.)

> **The EU AI Act enforcement deadline for high-risk AI systems is August 2, 2026. Non-compliance penalties reach EUR 35 million or 7% of global annual turnover, whichever is higher.**
>
> *Source: European Parliament, Regulation (EU) 2024/1689, Articles 5, 101–102. eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689*

The AWM™–Compliance Mapping

The AWM™ lifecycle maps directly to the EU AI Act's compliance requirements for high-risk AI systems. Acquisition addresses risk classification and vendor documentation requirements. Onboarding addresses technical documentation and access control requirements. Management addresses logging, human oversight, and ongoing risk management requirements. Offboarding addresses data governance and log retention requirements.

Organizations that implement AWM™ as described in this paper will, as a direct consequence, satisfy the majority of their EU AI Act compliance obligations for agentic AI deployed in high-risk categories. This is not coincidental. AWM™ is designed around the same governance logic that the regulatory framework codifies: know what agents you are operating, define and enforce their scope, maintain comprehensive behavioral records, and ensure appropriate human oversight.

SECTION 06

Building an AWM™ Program

AWM™ Maturity Self-Assessment

Before designing an AWM™ implementation, organizations should assess their current state across six governance dimensions: agent inventory, agent identity, scope definition, audit trail, oversight architecture, and offboarding process. Each dimension is scored on a four-level scale from 0 (no capability) to 3 (fully managed), producing a composite AWM™ Maturity Rating out of 18. Organizations scoring 0–6 are Pre-AWM and should treat implementation as immediate action. Scores of 7–11 indicate AWM Emerging; 12–15, AWM Defined; and 16–18, AWM Managed.

Most organizations conducting this assessment for the first time will score in the Pre-AWM or Emerging range. The value is not in the score itself but in identifying which dimensions represent the highest-priority gaps. A complete self-assessment worksheet is provided in the appendix.

Implementation Priorities

For most organizations, the assessment will reveal multiple gaps. Implementation should be sequenced to address the highest-risk exposures first.

1. **Agent Inventory and Classification**
   You cannot manage what you do not know exists. The immediate goal should be a complete registry of agents operating in the environment, with a preliminary risk classification applied to each. This is the prerequisite for every other AWM™ initiative.

2. **Identity Infrastructure**
   No other AWM™ capability functions without agent identity. Evaluate and deploy an agent credentialing solution before expanding agent deployments further. Deploying additional agents without an identity layer compounds the governance debt that will eventually need to be addressed.

3. **Scope Documentation and Enforcement for High-Risk Agents**
   Agents operating in high-risk classifications, including those touching HR, financial, or customer-facing processes, should have documented and technically enforced scope definitions before the August 2026 EU AI Act enforcement deadline.

4 **Audit Trail Implementation**

Every agent action should be logged in a manner that supports accountability attribution and regulatory compliance, regardless of where the organization sits on other AWM™ dimensions. This is the minimum viable compliance posture.

5 **Oversight Architecture Design**

Human oversight processes should be redesigned to function at the actual scale of agent deployment, with automation handling the volume that cannot be reviewed individually and human judgment applied to the escalations that genuinely require it.

Common Failure Modes

Treating AWM™ as an IT project rather than an organizational discipline

AWM™ requires policy decisions, oversight design, and governance accountability that belong to the business, not to IT. Technology implements AWM™; it does not define it. Organizations that delegate AWM™ entirely to IT will find that the governance questions, scope definition, oversight architecture, risk classification, remain unanswered.

Attempting to solve the governance problem with monitoring alone

Monitoring detects problems. Scope definition and access controls prevent them. A surveillance-heavy AWM™ program that lacks proper credentialing and access management is addressing symptoms rather than the underlying structural gap.

Designing for current scale rather than projected scale

AWM™ infrastructure adequate for ten agents will not be adequate for one hundred. Design for the agent footprint that is coming in two to three years, not the one that exists today.

Neglecting offboarding

Every AWM™ program should have an offboarding protocol before it has deployed a single agent. The failure to build offboarding into the program from the beginning is the organizational equivalent of hiring without an offboarding process: a problem that seems manageable until, suddenly, it is not.

Conclusion

The enterprise workforce has already changed. AI agents are not coming. They are here, operating inside enterprise environments, taking actions that affect employees, customers, and the organizations they represent. The governance frameworks designed for human workforces, and later extended for contingent workforces, were not built for this new workforce tier.

Agentic Workforce Management™ provides a framework for closing that gap. It applies the governance logic that organizations have developed over decades, for human workforces first, then for contingent and extended workforces, to the distinct requirements of autonomous AI agents: verifiable identity, defined and enforced scope, continuous behavioral monitoring, and systematic lifecycle management from acquisition through offboarding.

The organizations that build AWM™ programs now are not simply managing compliance risk. They are developing an organizational competency that will define their capacity to operate in an agentic environment. The frameworks, the processes, the infrastructure, these take time to build and time to mature. The organizations that start now will have a meaningful head start on those that wait for a forcing event.

The workforce already includes agents. The agentic era is not a future state to be planned for. It is the present state to be managed. The question is whether your organization manages it accordingly.

ABOUT THE AUTHOR

Michael Beygelman is Co-Founder and CEO of Insygna Corporation. A three-time founder with multiple exits, he previously built Claro Analytics (acquired 2022), led Adecco's global RPO and MSP businesses to over $1B in spend under management, and won the 2024 HR Tech Product of the Year for agentic AI. He writes The AI Workforce Chronicles, a newsletter on trust, accountability, and identity in the AI workforce.

ABOUT INSYGNA

Insygna Corporation provides trust infrastructure for the AI workforce. Its platform delivers verifiable identity, credentialing, trust scoring, and lifecycle management for autonomous AI agents in enterprise environments, enabling organizations to implement Agentic Workforce Management™ programs on a technical foundation designed for the demands of agentic AI. Insygna is headquartered in Merrimack, New Hampshire.

insygna.ai

APPENDIX

AWM™ Maturity Self-Assessment Worksheet

For each dimension, identify the description that most accurately reflects your organization's current state and record the corresponding score (0 to 3). Sum the scores to produce an AWM™ Maturity Rating. The maximum score is 18.

Governance Dimension	Maturity Levels, Score the level that best describes your current state	Score 0–3
Agent Inventory	0: No registry. Deployments are ad-hoc. 1: Leaders know agents are deployed but cannot enumerate them. 2: Registry exists but may be incomplete. **3: Complete, current registry with risk classifications.**	___
Agent Identity	0: No persistent identity. Agents share credentials. 1: Identity discussed but not implemented. 2: Some agents have distinct credentials; others share. **3: All agents have unique, verifiable identity credentials.**	___
Scope Definition	0: No documented scope. Agents operate without boundaries. 1: Scope in policy documents but not technically enforced. 2: Scope documented and partially enforced. **3: Scope encoded in credentials, enforced at runtime.**	___
Audit Trail	0: No agent action logging. Actions are unattributable. 1: System-level logs exist but not tied to agent identity. 2: Actions logged; attribution possible with manual effort. **3: Immutable, agent-attributed audit trail for all actions.**	___
Oversight Architecture	0: No escalation design. Human review is ad-hoc. 1: Human review expected but not systematically designed. 2: Escalation conditions defined; not fully automated. **3: Tiered oversight enforced technically, reviewed periodically.**	___
Offboarding Process	0: No offboarding process. Access persists after engagement. 1: Aware of risk but no formal process in place. 2: Checklist exists; execution inconsistent. **3: Systematic offboarding with revocation and archiving.**	___
Scoring guide: 0–6 = Pre-AWM (immediate action required) 7–11 = AWM Emerging 12–15 = AWM Defined 16–18 = AWM Managed		**Total: ___/18**