KPMG

Agentic AI Gateway

Data and Access Governance as an Enterprise Gateway for Agentic AI

Olle Jonsson
Jean El Khoury
Benjamin Gregersen

April 2026

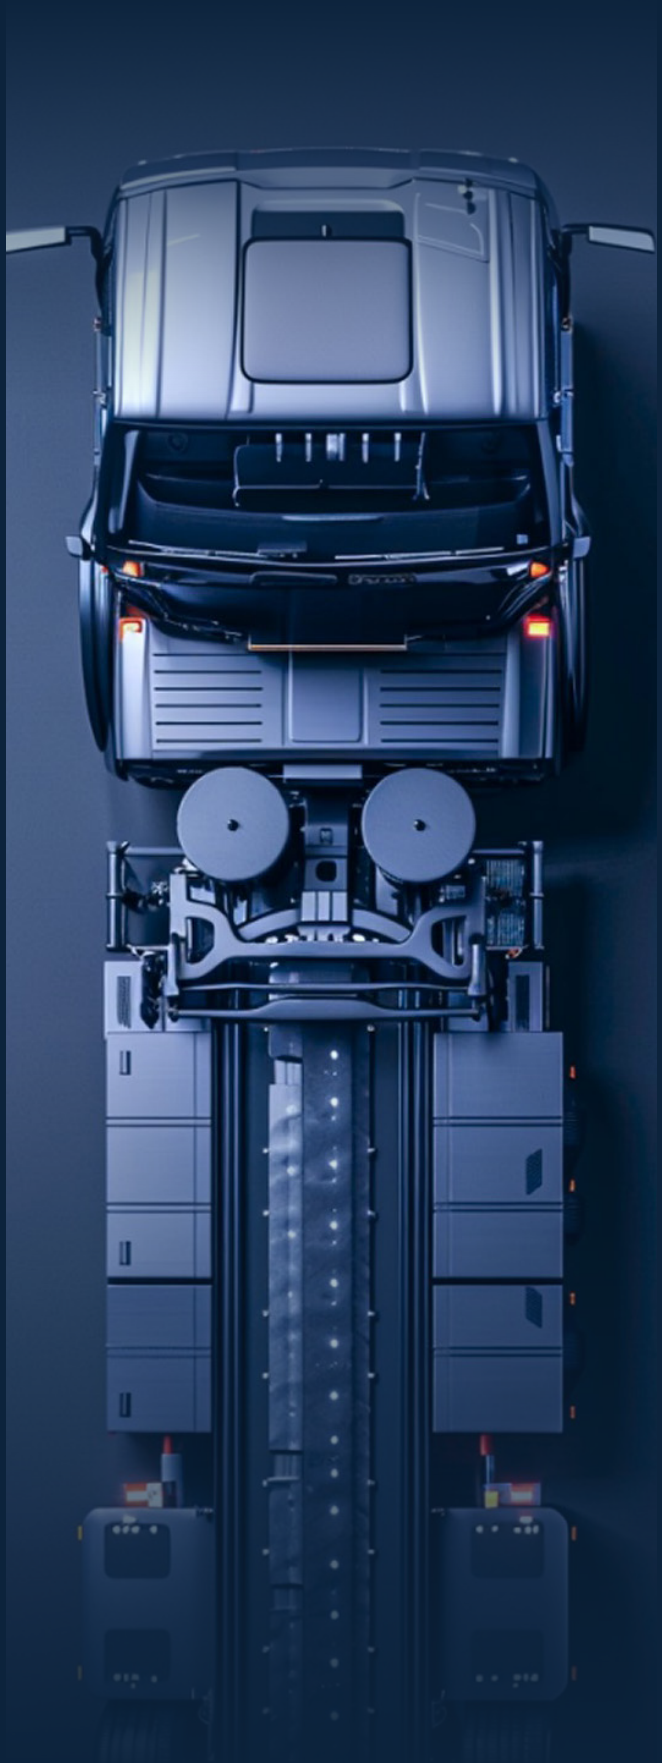

Table of Contents

Executive Summary

- Automotive is entering an always-on era, with AI agents rapidly expanding across the enterprise.
- Traditional IAM approaches were designed for human users and simply cannot keep up with the rising number of identities that are autonomous, ephemeral, and continuously active.
- Ungoverned AI agents create growing exposure across the organization, including security, compliance, and operations.
- The frontrunners will be those that govern AI agents through Clarity, Control, and Confidence.
- Leaders should act now to build visibility, enforce lifecycle control, and scale AI with confidence.

The Growing Complexity of Identity Governance in Automotive Industry

Automotive is becoming an increasingly software driven and connected industry. Vehicles, factories, and enterprise systems exchange data continuously, which expands the digital ecosystem and thus, the number of nonhuman identities (NHI) in daily use. NHIs are digital entities that independently access systems and perform tasks across the business, such as running routine operations or exchanging data. NHIs include applications, automated processes, system integrations, cloud services, connected devices and vehicles, as well as AI agents. For simplicity, this paper will focus on AI agents, however, the insights are equally applicable and relevant for other types of NHIs mentioned above. Just like human identities, NHI uses credentials (e.g., digital keys) to authenticate themselves and access organizational resources such as documents, systems, or operational data. These credentials ensure that NHIs, such as AI agents, can only reach the information they are allowed to see or use.

As automotive companies are under competitive pressure, they seek to embed AI deeper into their core processes to operate faster, at greater scale, and with fewer manual handovers. Advances in cloud platforms, large language models, and enterprise automation now allow AI agents to act increasingly as digital co workers, augmenting experts in monitoring information, making recommendations, and triggering actions continuously rather than on demand. This shift will only accelerate as automotive companies advance toward always-on, data-driven operations.

Today, AI agents are mainly used to support customer interactions, monitor and optimize production and supply chains, automate IT and engineering workflows, and continuously conduct compliance, quality, and risk checks across complex, connected ecosystems. While organizations have focused on managing human identities for years, they now find themselves challenged by the rapid increase of NHIs. According to Entro Security's NHI & Secrets Risk Report (H1 2025), non-human identities now outnumber human identities by approximately 144:1. This rapid growth vastly increases complexity for governance and assurance. Therefore, organizations need to adapt and build a technical backbone that scales with the growing number of NHIs to capture the innovative force Agentic AI fully without compromising on security aspects.

Why Traditional Approaches to Managing AI Agents Fail

Traditional Identity and Access Management (IAM) approaches were designed for human users: employees who join an organization, change roles occasionally, and eventually leave. These models rely on manual approvals, static roles, and periodic access reviews, assuming relatively stable identities and predictable behavior. NHIs, such as AI agents, operate very differently: they are created based on demand and scaled in large numbers, act around the clock, interact directly with other systems, and often evolve or disappear without clear ownership. As a result, traditional IAM cannot keep pace, resulting in access remaining active longer than needed, limited visibility into what agents do, and unclear accountability. This creates blind spots where risk accumulates, audits become harder to explain, and operational friction increases, highlighting the need for new governance approaches that reflect the speed, scale, and autonomy of NHIs. In fact, Saviynt's Identity Security for AI report (2026) highlights that 91 % of enterprises lack sufficient visibility into their AI and non-human identities.

New Types of Identities – new challenges

Organizations are eager to harness the innovative potential of AI agents to drive efficiency, automation, and new sources of value. Yet this ambition also creates tension between rapid experimentation and the need to safeguard sensitive data, systems, and trust. Without the right controls, AI agents can accumulate security, compliance, and operational risks that differ fundamentally from those of human users. Drawing on selected themes from the OWASP Non-Human Identities Top 10 (2025), the examples below highlight some of the most critical risks emerging from the rapid growth of NHIs.

Uncontrolled AI Identities Increasing Exposure

AI agents are often created quickly to support pilots, automations, or temporary initiatives, but are rarely retired with the same discipline as human users. When agents or their credentials remain active after their business purpose ends, they create long lived and invisible access paths into critical systems. Over time, these "orphaned" agents accumulate and significantly increase security and audit risk, as organizations lose track of which agents exist, which still have access and why.

Credential Leaks Enabling Undetected Breaches

AI agents rely on digital secrets such as API keys, tokens, and certificates to authenticate and act. These secrets are frequently exposed through source code repositories, configuration files, or collaboration tools, often unintentionally. Once leaked, they allow silent and automated access to systems without triggering traditional security alarms, making secret leakage one of the most common and impactful entry points for misuse involving AI agents.

Third Party Weaknesses Extending Enterprise Risk

Many AI agents depend on third party platforms, SaaS services, or external tools to function effectively. When these third party identities are compromised or poorly secured, attackers can inherit the trust and access granted to the agent. This extends risk beyond organizational boundaries and makes accountability more complex, as incidents may originate outside direct control while still impacting core business systems.

Insufficient Authentication Eroding Trust

AI agents are often authenticated using outdated, weak, or inconsistent mechanisms that were never designed for highly autonomous, always on actors. Insecure authentication reduces confidence that actions can be reliably attributed to a specific agent and increases the likelihood that credentials can be replayed or misused. For organizations, this undermines both security and the ability to explain agent behavior during audits or investigations.

Over Privileged Agents Magnifying Incident Impact

To avoid interruptions, AI agents are frequently granted broad or excessive access across systems and data. While convenient, this practice dramatically increases potential impact if an agent is compromised or behaves unexpectedly. Over privileged agents can move laterally across environments, access sensitive data beyond their intended scope, and amplify the business consequences of a single failure or misuse.

Getting in Front of the Wave: Clarity, Control, and Confidence over Your AI Identities

Today, most organizations already have many different types of AI agents operating across the organization: business chatbots embedded in SaaS tools, workflow automations in cloud platforms, copilots in productivity suites, machine learning agents in factories, and custom models built for more specific use cases to name just a few. These agents sit in different environments, rely on a range of vendors, and have inconsistent levels of oversight. As the digital ecosystem grows, singular or tool specific approaches simply cannot scale. What organizations need instead is an enterprise wide, vendor agnostic framework that governs all AI agents the same way, regardless of where they run or who provides them.

The three C Framework

This is where the three C's clarity, control, and confidence become essential. The Three C framework is a proprietary methodology developed by KPMG specifically for AI identity governance.

- *Clarity* means knowing exactly which AI agents exist, what they do, and who is responsible for them.
- *Control* means ensuring agents receive only the access they need, only for as long as they need it, and that their actions can be verified.
- *Confidence* means being able to trust the system because activity is monitored, evidence is consistent, and audits can be answered quickly.

By applying these three principles, organizations can reduce complexity, improve operational efficiency, and create a safer environment for adopting agentic AI. The three C is a practical framework that helps business leaders bring order to a fast growing digital workforce and govern AI agents consistently across the whole enterprise.

	Clarity	**Control**	**Confidence**
Definition	Know every agent operating in the organization, its purpose and lifecycle status.	Allow access only when needed and only for the time the need exists.	Know what agents are doing with evidence that stands up to audit and scrutiny.
Good practice	Enterprise inventory with owner, purpose, access scope, and life-cycle state and a consolidated view across environments.	Deny-by-default; just-in-time access; action verification for sensitive operations; automated onboarding / offboarding and periodic attestation.	Tamper-resistant audit trails; continuous monitoring; ready-to-use information to answer incidents, audits, and regulators.
Business benefits	Faster onboarding through reuse of approved patterns; spend visibility and tool rationalization.	Reduced attack surface and blast radius; improved time-to-market through efficient onboarding and offboarding.	Audit-ready posture; confidence to scale pilots to production with governed evidence.

Clarity

As automotive organizations accelerate toward software defined vehicles, connected factories, and data driven mobility services, they already operate thousands of different AI agents, telematics modules, over the air (OTA) update services, manufacturing robots, and diagnostic tools. This creates an extremely complex ecosystem, making manual governance impossible to sustain.

Achieving *Clarity* means establishing an enterprise wide, vendor agnostic inventory of all these agents regardless of whether they live in the vehicle, the cloud, or on premises. It includes knowing who owns each agent, what its purpose is, where it runs, what data it can access, and how it behaves. For example, a connected car platform may involve hundreds of microservices, each with its own identity. Without clarity, dormant or "orphaned" identities remain active, creating blind spots that attackers can exploit.

In practice, clarity comes from consolidating identity data across all platforms. This unified view helps teams recognize overlaps, reuse proven identity patterns, and accelerate time to market for new AI powered services such as predictive maintenance, smart charging, or fleet management applications. Instead of every team reinventing its own approach, an enterprise inventory provides a single source of truth for innovation and governance.

Control

AI agents interact constantly across domains. If these machine identities have standing, long lived access across these systems, the blast radius of a compromise becomes enormous. This is exactly what the industry is increasingly seeing, as static credentials or unmanaged API keys contribute to breaches across modern mobility platforms.

Control means shifting from open ended, persistent access to just in time, purpose bound permissions. For example, a software update service should only receive access when an OTA package is deployed, not 24/7. A factory robot should authenticate with short lived credentials tied to its shift cycle. A diagnostic tool used by an external dealer should receive access for the duration of a repair session and then lose it immediately afterward.

Trends in the market highlight how organizations are moving toward Zero Trust principles, where every identity human or non human is continuously verified and never implicitly trusted. For example, if a service account suddenly starts interacting with systems outside its usual pattern (e.g., a telematics API accessing manufacturing data), AI driven monitoring can flag abnormal behavior and enforce additional control steps or revoke access completely.

Establishing Control also requires automated onboarding and offboarding for agents based on clear principles, ensuring that identities are created and removed through policy not

emails, spreadsheets, or manual "clean up" cycles. This reduces both cyber risk and operational friction, enabling engineering and DevOps teams to introduce new services or train new AI agents with far shorter lead times without going to manual approval workflows.

Confidence

The challenge is not only knowing which AI identities exist (clarity) or limiting what they can do (control) it is being able to prove, at any moment, that they are behaving as intended. Confidence means having a governance and monitoring model where AI agents' actions are continuously observed, recorded, and explainable.

Practically, confidence is built through tamper resistant audit trails, continuous monitoring, and ready to use evidence that can answer incidents, audits, and regulator inquiries quickly without weeks of manual log stitching across vendors and platforms. For example, if an OTA update service deploys a package, confidence means you can reconstruct the full chain of events within minutes: authentication, authorization, package integrity checks, deployment actions, and post deployment validation mapped to the responsible owner and the approved change record.

Ultimately, when evidence is standardized and continuously collected, organizations avoid the time-consuming audits, shorten investigations, and can demonstrate governance maturity to customers, partners, insurers, and regulators. In other words, confidence is the capability to scale AI agents at high speed, without compromising on security, because oversight is embedded into the operating model rather than responding in case something goes wrong.

Implementing the Three Cs Through a Layered, Vendor Agnostic Architecture

To achieve **Clarity, Control and Confidence**, organizations need to build their architecture around three complementary layers an **AI agent ecosystem**, a control layer, and a **data layer**. Together, these layers form the foundation for achieving Clarity, Control, and Confidence.

AI Agent Ecosystem

AI agents are no longer confined to a single platform or provider; they run across cloud infrastructure, SaaS applications, and custom technology stacks, and are delivered by a mix of hyperscalers, software vendors, and internal development teams. Governance must therefore be inherently decoupled from technology choices and vendor dependencies. This ecosystem layer ensures that all AI agents regardless of where they run or who provides them are recognized as enterprise identities that fall under the same governance principles.

Control Layer

The control layer focuses on managing access and behavior of the AI agents. This layer applies identity and access management (IAM) principles, defining how identities are created, how permissions are granted and revoked, and how actions are verified. The control layer is what turns policies into enforceable rules, ensuring AI agents operate with purpose bound, time limited, and verifiable access.

Data Layer

The data layer forms the foundation of the entire model by governing data classification, usage, and quality. It provides transparency into what data AI agents consume, how sensitive information is handled, and how data flows across environments. This is critical not only for protecting sensitive information but also for ensuring that decisions made by AI agents can be understood, justified, and reviewed. Without a reliable data layer, neither effective access control nor the expected business value from Agentic AI is possible to achieve.

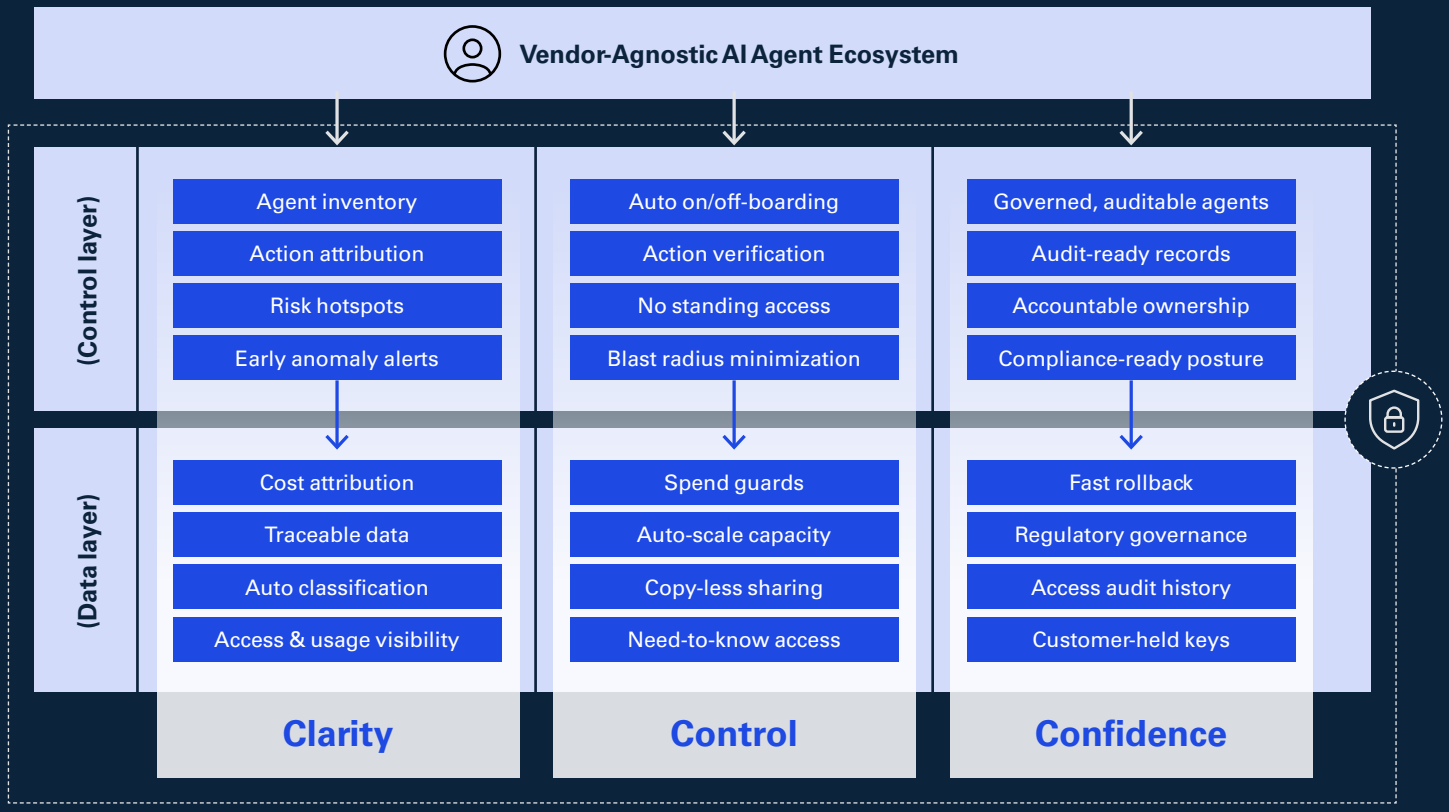

Only when all three layers are in place organizations can reliably achieve Clarity over what AI agents exist, Control over what they are allowed to do, and Confidence that their behavior is governed, auditable, and aligned with internal policies and regulatory requirements.

Case Study: Streamlining contract management with AI agents in Automotive

A proven example of the three C framework in action comes from an automotive organization looking to modernize its contract renewal and cost optimization process. The business goal was straightforward: identify all contracts expiring within the next 120 days, uncover saving opportunities, and prepare clear renewal recommendations without relying on time-consuming manual analysis. To achieve this, the organization deployed a set of specialized AI agents, each focused on a specific task such as detecting upcoming expirations, analyzing supplier terms and pricing, and generating negotiation recommendations for decision makers.

These agents were implemented within a vendor agnostic AI agent ecosystem, allowing them to operate seamlessly across different platforms and systems. Contract, supplier, operational, and financial data were spread across ERP systems, pricing tools, and contract repositories, making platform specific solutions impractical. By decoupling governance from vendors and technologies, the organization ensured consistent oversight regardless of where the agents ran or who provided them. This approach also created a scalable foundation for future AI enabled use cases beyond contract management.

The control layer ensured consistent and scalable governance of the AI agents themselves. Each agent was treated as a non human identity with a defined purpose, ownership, and scope of access. Identity and access management principles were applied to manage onboarding, permissions, and monitoring, while agent actions were continuously logged and attributed. This made it possible to trace which agent accessed which systems and data, and why.

The data layer proved critical in making the results both reliable and actionable. AI agents were granted controlled access to contract data, operational data, supplier data, and ERP and pricing information based on clearly defined classifications and usage rules. Data traceability, quality controls, and fit for purpose access ensured that insights were grounded in trusted information while preventing unnecessary exposure of sensitive data. As a result, the organization could quantify savings opportunities with confidence rather than relying on assumptions or partial views of the data.

Together, these layers enabled the practical application of the three Cs. Clarity was achieved through transparency over the agents involved, their responsibilities, and the data they relied on, turning contract renewals into a structured, repeatable process. Control was enforced by ensuring agents operated with purpose bound and time limited access, minimizing risk without slowing analysis. Confidence emerged from traceable data usage and monitored agent activity, allowing outputs to be trusted, reused, and defended in internal reviews.

The business impact was immediate and measurable. In less than ten minutes, the governed AI agents identified more than 30 contracts fulfilling the defined criteria, identified approximately 11 MSEK in potential savings, and highlighted key negotiation levers. Manual effort associated with data extraction and consolidation was eliminated, full financial transparency across contracts was established, and the client strengthened its negotiation position with suppliers. Most importantly, proactive visibility ensured that no contract expirations or price increases were missed. This example demonstrates how the three C framework enables organizations to scale AI agents safely delivering speed, financial value, and confidence without compromising governance or compliance.

Conclusion: A structured Plan to Drive Innovation at Scale

To get started with a scalable, vendor-agnostic approach to governing AI agents, organizations must first establish strategic clarity such as agreeing on priorities for AI adoption, and risk appetite. These high-level choices should then be translated into practical implications for data usage, access rights, accountability, and lifecycle management of AI agents in day-to-day operations. Rather than building from scratch, organizations should deliberately build on existing data management, IAM, and governance structures, using them as a foundation to incrementally evolve toward a defined target state. This step-by-step transition enables scale and consistency while maintaining control as AI agent adoption accelerates across the enterprise.

Key actions in this context include:

---

1. Define the enterprise meaning of the three Cs

Clearly articulate what control, compliance, and consistency mean for your organization and translate them into guiding principles that apply across all functions and technology domains – rather than addressing risks through fragmented, solution specific initiatives.

---

2. Assess and leverage existing in house capabilities

Evaluate current data management and identity and access management capabilities to determine what controls, processes, and governance mechanisms are already in place and can be extended to AI agents and other non human identities.

---

3. Establish a single source of truth for AI agents and NHIs

Build a central repository that provides transparency over existing AI agents and non human identities across the enterprise. Where needed, this may be supported by a targeted vendor selection process to ensure scalability and integration with existing platforms.

---

4. Implement standardized onboarding and offboarding processes

Define clear lifecycle processes for introducing, managing, and retiring AI agents to ensure continuous visibility, accountability, and risk management as adoption accelerates.

Building upon this foundation, organizations can extend the technical backbone step by step, to govern the growing number of AI agents in an efficient and scalable way to succeed in today´s AI race within the automotive industry.

Take the Lead in Identity-Driven Transformation

Governing machine identities requires a strategy built on architecture, operations, and AI-driven intelligence. Contact us to build and execute scalable, AI-powered identity governance.

Olle Jonsson
Director, Cyber Security

+46 790 658 878
olle.jonsson@kpmg.se

Jean El Khoury
Associate, Cyber Security

+46 8 723 91 00
jean.el.khoury@kpmg.se

Benjamin Gregersen
Associate, CIO Advisory

+46 707 232 669
benjamin.gregersen@kpmg.se

Daniel Szirányi
Global Commercial Vehicles CoE Lead
Partner, MC Core

+46 736 990 547
daniel.szirayani@kpmg.se

KPMG

Visiting address:
Vasagatan 16, Stockholm
Postal address:
P.O. Box 382
SE-101 27 Stockholm

Tel: +46 8 723 91 00
E-mail: info@kpmg.se

kpmg.se