unosecur

WHITEPAPER

Identity Security in the Automotive Industry

Why is identity the common thread through every major automotive cyberattack, and what a modern defense looks like.

Table of contents

Executive summary

The automotive industry is undergoing its most significant transformation since the introduction of mass production. Vehicles are now software platforms. Factories are cloud-connected. Supplier ecosystems span hundreds of organizations sharing systems, data, and credentials. And AI agents are beginning to operate autonomously within these environments.

This transformation has made the automotive sector one of the most targeted industries in cybersecurity. In 2025 alone, Upstream Security recorded 494 publicly reported incidents, more than double the prior year. Ransomware accounted for 44 percent of attacks. The JLR incident resulted in a five-week production shutdown, direct losses estimated in the hundreds of millions, and $4.69 billion in emergency financing. The CDK Global breach brought 15,000 dealerships to a halt after attackers gained access to its dealer management platform, disrupting operations across the US auto retail sector and causing estimated losses exceeding $1 billion.

The pattern across every major incident is consistent. Identity is the attack surface. Stolen credentials, misconfigured IAM roles, leaked API tokens, and overprivileged service accounts are not supporting actors in these attacks. They are the mechanism by which adversaries move from initial access to full operational disruption.

494	2x	44%	92%
Publicly reported incidents in 2025	Year-on-year growth in attacks	of incidents involved ransomware	of attacks conducted remotely

This paper maps the identity threat landscape specific to the automotive industry, examines the incidents that define it, and makes the case for an identity-first security posture. It also introduces the emerging frontier: AI agents, which are now operating as autonomous identities within automotive environments, often without the governance frameworks to their risk profile demands.

Section 1: The automotive identity landscape

Most enterprises have complex identity environments. Automotive enterprises have an extreme one. Three distinct identity types coexist across every major OEM, supplier, and fleet operator, and each carries a different risk profile.

Human identities

Employees, contractors, and supplier personnel who directly access systems. In automotive, this category is complicated by scale: a single OEM may have relationships with 600 or more Tier-1 and Tier-2 suppliers, each with its own access credentials and entitlements. These identities are the most familiar to security teams, but third-party human identities operating across supplier boundaries remain poorly governed in most organizations.

Non-human identities (NHIs)

APIs, service accounts, tokens, certificates, and the ECUs embedded in connected vehicles. Modern vehicles carry over 100 ECUs, each representing a non-human identity with credentials and communication paths. The ratio of non-human to human identities is already at a concerning 144:1. These identities are rarely rotated, often untracked, and frequently overprivileged. They are also the fastest-growing attack vector in the sector.

AI agents

The newest and least-governed identity class. AI agents in automotive environments manage fleet telemetry, orchestrate OTA firmware updates, power predictive maintenance systems, and interface with vehicle APIs. Unlike a service account that executes a defined task, an AI agent makes decisions, chains actions, and operates across multiple systems. It requires an identity with broad permissions, and in most current deployments, it receives them without the controls that its autonomous behavior demands.

The automotive industry does not have a vehicle security problem. It has an identity governance problem at a planetary scale.

The convergence of these three identity types across cloud infrastructure, SaaS platforms, on-premise factory systems, and operational technology create a vast, fragmented identity perimeter that no organization in this sector has fully mapped, let alone secured.

Section 2: How identity is being attacked

The following incidents represent recurring patterns of identity exploitation across the automotive sector. Each is documented, attributed, and instructive.

1. **2025 | Jaguar Land Rover**
   Pattern: Supply chain + human credential abuse. Kill-chain ransomware attack shuts global production for 40 days

   Attackers gained initial access through spear-phishing and credential theft, then used valid accounts to persist, escalate privileges via overpermissive IAM roles, and move laterally through IT and OT environments using legitimate remote services. The result was a five-week worldwide production shutdown, losses estimated in the hundreds of millions, and $4.69 billion in emergency financing. The MITRE ATT&CK mapping of this incident shows identity abuse at every stage of the kill chain, from initial access through to impact.

2. **2025 | Volkswagen / 8Base**
   Pattern: Third-party identity entry point. Ransomware group claims theft of sensitive internal data.

   The 8Base group, believed to be an offshoot of the Phobos ransomware family, claimed responsibility for a breach that allegedly exposed invoices, employment contracts, personnel records, and confidentiality agreements. Volkswagen confirmed an incident had occurred while stating no IT systems were impacted, consistent with a compromise originating through a third-party supplier or subsidiary. Initial access is consistent with the group's known approach: phishing campaigns or credentials purchased on the dark web.

3. **2025 | Tata Motors**
   Pattern: NHI / secrets sprawl. Hardcoded cloud keys expose massive enterprise data lake

   Hardcoded AWS access keys were embedded in publicly accessible application code powering Tata Motors' E-Dukaan portal and FleetEdge platform. Security researchers found that these credentials acted as master keys to the company's cloud infrastructure, granting access to over 70 TB of sensitive data, including customer PII, financial records, fleet telemetry, and internal dashboards. Additional exposed API tokens enabled access to vehicle tracking systems, while a "trusted token" mechanism allowed passwordless admin-level access to internal analytics platforms.

No sophisticated exploit was required. No perimeter was breached. A handful of unmanaged non-human identities, left exposed in code, unlocked the entire ecosystem.

4. **2023 | Volkswagen / Cariad**
   Pattern: Cloud IAM misconfiguration exposes the location data of 800,000 EV owners

   VW's Cariad software division misconfigured AWS access controls, leaving precise GPS location data and personal information for approximately 800,000 electric vehicle owners were exposed for several months. The root cause was an improperly scoped cloud IAM policy. The breach required neither credential theft nor exploitation of a software vulnerability. Misconfigured machine identity permissions were sufficient.

5. **2024 | CDK Global**
   Pattern: Human credential + lateral movement. A single set of phished credentials halt 15,000 dealerships

   BlackSuit ransomware operators gained access to CDK Global's dealer management platform through a phishing attack targeting employee credentials. Because CDK Global serves as a centralized identity and operations platform for dealerships across the United States, lateral movement from a single compromised account resulted in an operational shutdown across 15,000 sites. Estimated industry losses exceeded $1 billion.

The five identity attack patterns

Across these and other incidents, five patterns recur:

- **Supply chain identity backdoor:** supplier or contractor credentials used to access OEM systems
- **Human credential theft:** phishing, vishing, or dark web purchase of a valid employee accounts
- **NHI and secrets sprawl:** unmanaged tokens, API keys, or certificates left exposed in code, firmware, or cloud configuration.
- **Cloud IAM misconfiguration:** overpermissive policies granting access that was never intended
- **Lateral movement on valid accounts:** using legitimate credentials to pivot across IT, OT, and cloud environments without triggering conventional detection

In every case, the infrastructure was not broken. The access was.

These five patterns define the current threat landscape. But the environment is not static. As automotive enterprises deploy AI agents across fleet operations, vehicle software, and supply chain systems, each of these patterns is being amplified.

An AI agent is simultaneously a human-equivalent actor with credentials, a non-human identity with API access, and an autonomous decision-maker operating across multiple systems. It inherits every vulnerability from both identity categories while introducing risks that neither was designed to address. That convergence is what makes AI agent identity governance the most urgent emerging challenge in automotive security.

Section 3: The emerging frontier: AI agents as identities

AI-Powered psychological manipulation at scale

AI agents are already operating in automotive environments: managing fleet telemetry, orchestrating OTA updates, and coordinating V2X communications. In most deployments, they hold broad permissions because restricting them slows the AI adoption that the business is prioritizing. That tradeoff is creating a governance gap that threat actors are beginning to understand and exploit.

Four automotive-specific risk scenarios

1. **OTA update agent compromised:** an agent responsible for firmware distribution is hijacked via prompt injection or credential theft. The attacker uses the agent's existing permissions to push malicious firmware to an entire vehicle fleet. No individual login is required.

2. **Fleet analytics agent credential leak:** an agent accessing vehicle telemetry and driver data has its API credentials harvested via malware. The attacker gains persistent access to live vehicle location data and personal information for millions of users.

3. **Customer support AI prompt-injected:** a customer-facing AI agent is manipulated through a crafted user input to retrieve and expose vehicle access tokens. This is not input manipulation in the traditional sense. It is identity hijacking through instruction abuse.

4. V2X coordination agent spoofed: an agent coordinating vehicle-to-infrastructure communications is compromised and used to manipulate traffic signaling or feed false data to autonomous vehicle systems. The agent's trusted identity is the attack vector.

The structural risks behind these scenarios

- Confused deputy: The agent has legitimate permission to access a system; the attacker tricks the agent into using that permission on their behalf.
- Cross-agent propagation: In multi-agent architectures, one compromised agent can pass its credentials or manipulate downstream agents, spreading access across the ecosystem.
- Orphaned credentials: Agents created for short-lived tasks often have credentials that persist long after the task ends.
- Audit gap: When an action is taken by an agent acting on behalf of a user acting through an interface, attribution becomes genuinely difficult.

AI adoption in enterprise is outpacing identity governance. Less than a third of organizations enforce security controls when doing so would slow AI deployment. In automotive, where agents are interfacing directly with vehicle systems and supply chain infrastructure, this gap carries consequences that extend well beyond data loss.

Section 4: Why this is hard to solve

Understanding the attack patterns is necessary but not sufficient. CISOs in the automotive sector face structural conditions that make identity security genuinely difficult, independent of investment and intent.

Supplier identity perimeter sprawl

A major OEM works with 600+ Tier-1 and Tier-2 suppliers. Each supplier operates its own identity environment with its own access policies, credential standards, and offboarding processes. Federated access across this ecosystem is rarely implemented consistently. Every supplier relationship is a potential identity boundary with uncertain controls on the other side.

NHI lifecycle neglect

Non-human identities are created faster than they are tracked. API keys embedded in vehicle firmware, service account tokens provisioned for a deployment pipeline, certificates issued for a supplier integration: each has a lifecycle that, in practice, is rarely managed. Credentials persist long after the systems or relationships they served have changed. Rotation is infrequent. Scope is rarely enforced at creation.

OT/IT convergence without identity controls

Factory floor systems, programmable logic controllers, and manufacturing execution systems were designed for isolated networks. They carry no native support for modern identity controls: no MFA, no role-based access policies, no audit logging in the format that security tools expect. As these systems connect to corporate networks and cloud environments, the identity perimeter expands into territory that security teams were not built to govern.

Secrets sprawl

Credentials embedded in source code, firmware images, CI/CD pipelines, and configuration files represent a largely invisible attack surface. Secret sprawl is the fastest-growing identity vulnerability category in automotive, and also the least visible. Most organizations do not have a complete inventory of where their secrets are stored, let alone which ones are still active.

Lack of unified identity visibility

Human identities, NHIs, and AI agents are governed by separate tools, separate teams, and separate policies. There is no single view of who has access to what, whether that access is appropriate, and where anomalous behavior is occurring. Without that visibility, identity risk cannot be measured or managed.

Regulatory fragmentation

UN Regulation R155, ISO/SAE 21434, the EU Cyber Resilience Act, and regional data protection frameworks impose overlapping but inconsistent identity and access requirements across markets. Compliance activities consume security team capacity without necessarily delivering the unified identity posture that the threat environment demands.

Section 5: What an identity-first defense looks like

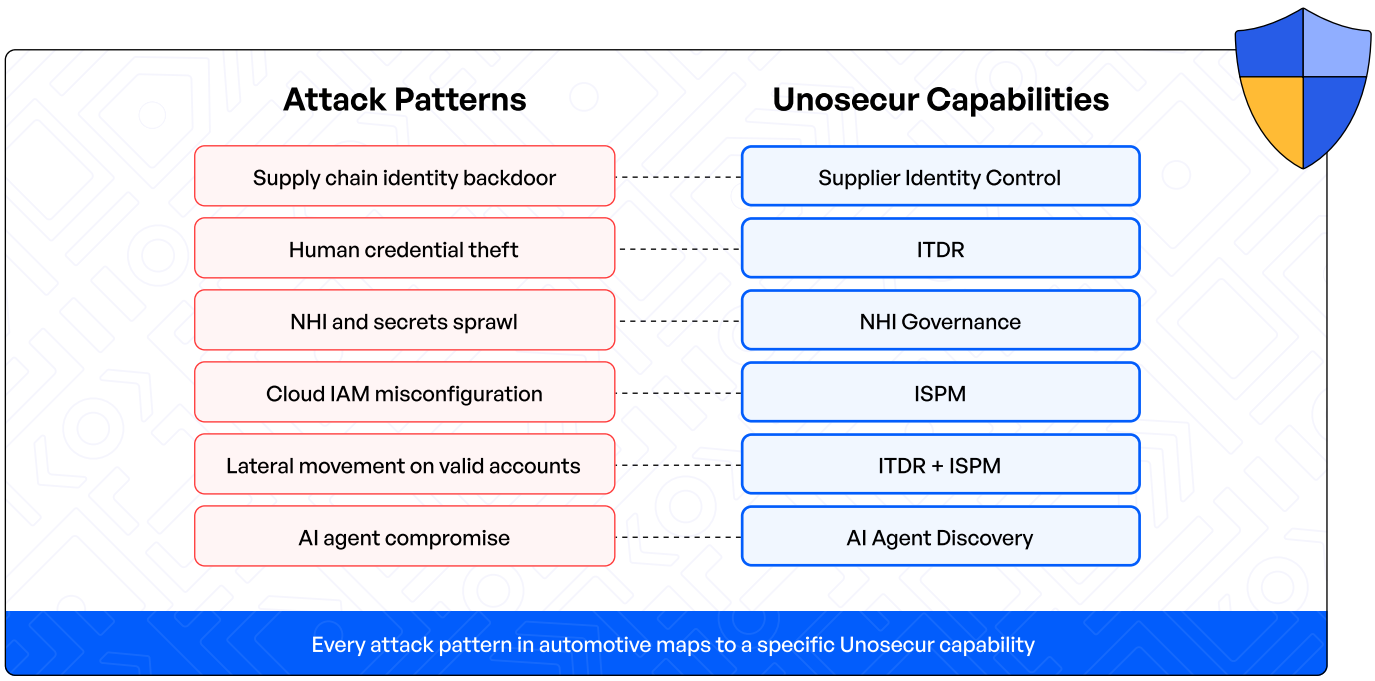

An effective response to the automotive identity threat landscape requires visibility and control across all three identity types: human, non-human, and AI agents. It requires that coverage to extend across cloud, SaaS, and on-premise environments simultaneously. And it requires the capability to detect anomalous identity behavior in real time, not after the fact.

Unosecur's Unified Identity Fabric is built to provide exactly that. The platform manages and governs human identities, NHIs, and AI agents within a single framework, designed for the complexity of enterprise environments.

ITDR (Identity Threat Detection and Response)

Continuously monitors identity behavior across environments, flagging anomalies such as credential use outside normal patterns, lateral movement on valid accounts, and privilege use inconsistent with the role. Directly addresses the kill-chain patterns in incidents involving JLR, CDK Global, and Nissan.

ISPM (Identity Security Posture Management)

Continuously surfaces IAM misconfigurations, overpermissive roles, orphaned accounts, and policy violations before they become exploitable. Addresses the cloud IAM misconfiguration pattern in Cariad and Mercedes incidents, as well as the persistent over-permissioning that enables lateral movement.

AI Agent Discovery

Discovers, inventories, and governs AI agent identities operating within the environment. Identifies agents with excessive permissions, tracks credential use by autonomous agents, and enforces least-privilege policies across multi-agent architectures. Built for the governance gap that automotive AI adoption is creating.

NHI Governance

Provides lifecycle management for non-human identities, including API keys, service account tokens, certificates, and embedded credentials. Tracks rotation status, flags unmanaged or long-lived credentials, and enforces scoping policies. Addresses secrets sprawl and NHI lifecycle neglect across the automotive environment.

Supplier Identity Control

Extends identity governance to third-party and supplier access, providing visibility into external identities within the OEM environment, enforcing least-privilege policies on cross-boundary access, and flagging anomalous supplier credential behavior. Addresses the supply chain identity backdoor pattern directly.

Unosecur provides the identity-first defense layer that breaks the kill chain before adversaries reach production systems and supply chain infrastructure.

Conclusion

Understanding the attack patterns is necessary but not sufficient. CISOs in the automotive sector face structural conditions that make identity security genuinely difficult, independent of investment and intent.

Supplier identity perimeter sprawl

The automotive industry's digital transformation has created conditions that threat actors are actively and systematically exploiting. The incidents of the past three years are not anomalies. They are the consequence of an identity perimeter that has expanded faster than the governance frameworks designed to protect it.

The JLR attack showed that identity abuse can shut down global manufacturing for

forty days. The CDK Global incident showed that a single phished credential can bring 15,000 operations to a standstill. The VW and Mercedes incidents showed that non-human identities, without lifecycle management, can expose millions of records without a single attacker having to authenticate directly.

AI agents are making this more urgent, not less. As automotive enterprises deploy autonomous agents across fleet management, vehicle software, and supply chain systems, they are creating a new class of identities that operate with broad permissions, make autonomous decisions, and are currently governed by almost nothing.

Three things every automotive CISO should do now:

- **Inventory every identity type across your environment:** human, non-human, and AI agent. If you cannot answer how many identities you have and what each one can access, you cannot begin to manage the risk.
- **Enforce least-privilege across all third-party and supplier access:** The supply chain identity backdoor is the dominant initial access vector in the automotive industry. It is also one of the most addressable, given the right governance tooling.
- **Establish AI agent identity governance before your agent deployment scales further:** The window between adoption and governance failure is short, and the consequences of closing it after an incident rather than before are significant.

The organizations that will be hit by the next wave of automotive cyberattacks will not be those with the most perimeter defenses. They will be those who have treated identity as the primary security control, consistently governed it across all three identity types, and built detection capabilities to act before lateral movement completes.