ENABLIS INSIGHTS // AGENTIC AI IN FLEET

Agentic AI in Fleet Management: Security, governance, & what good looks like

Agentic AI is already doing real work in fleet: scheduling, routing, maintenance and compliance. The question now is **how you run it well: what it's allowed to do on its own, how you keep it secure, and how you turn it into dependable, everyday delivery.** A practical read for fleet leaders on implementing agentic AI with confidence, covering security, governance, and what good looks like.

enablis.

// WHY THIS MATTERS

enablis.

The AI spectrum in fleet, today

Turning ambition into enterprise-level production is the real issue in fleet AI. 70% of UK transport and logistics leaders expect AI to transform fleet management, and more than eight in ten already use AI tools in some form (Microlise).

Yet only 5.6% run AI broadly, and most are still in pilots (Fleetio). Across every industry, just 17% of organisations have agents in production, even as 60% plan to within two years (Gartner). The distance between intent and execution is where the real work sits, and where the advantage is won.

70%
of UK transport leaders expect AI to transform fleet (Microlise)

5.6%
of fleets actually use AI broadly; most are still piloting (Fleetio)

17%
of organisations have deployed AI agents in production (Gartner)

From answering questions to doing the work

It's worth being precise, because "AI" now spans a wide spectrum. Predictive models forecast a breakdown, flag a risky driver or suggest a faster route. Assistants draft and summarise on request. Agents act: when a fault code fires, an agent can raise the work order, book the slot, order the part and tell the driver, with human involvement bedded in exactly where the business wants it. The security and governance questions differ sharply across that range.

Most fleets still sit at the predictive end of that spectrum, but the market is moving quickly towards the agentic one. **Three examples are already live:**

1. GEOTAB

Its June 2026 release turns a multi-step task into a conversation: a fleet manager can pull the data, spot the problem and act on it, all by chatting with ChatGPT, Claude or Microsoft Copilot. Central Transport is already using it to replace weeks of manual analysis with instant reporting.

2. MICROLISE TRUCONTROL

An AI assistant that answers an operator's compliance questions and points them to the specific action they need to take.

3. MICHELIN AI ASSISTANT

Now live in the UK, it lets a fleet manager ask questions of their own vehicle data in plain language rather than waiting on someone to build a report.

> **My conviction is simple and increasingly contrarian: AI is a productivity multiplier for people, not a replacement for them. I've put evidence behind it. Agentic AI embedded across MICHELIN Connected Fleet's software engineering lifecycle lifted delivery velocity 49% in six months. AI-led transformation across six commercial functions delivered seven-figure cost avoidance in a year - without a single role eliminated.**

Taimour ElSakka, CTO
MICHELIN Connected Fleet

More autonomy, more risk

As AI moves from advising to acting, the risk profile moves with it. A cyberattack on a connected fleet becomes a safety event, telematics data becomes a privacy liability. It's telling that Gartner expects more than 40% of agentic AI projects across all industries to be scrapped by the end of 2027, undone not by the technology but by unclear value and weak controls. The capability is already here; what makes the difference is the governance built around it.

// INDUSTRY ANGLE

enablis.

AutoTrader

Helping dealers advertise with AI

The same principle holds beyond fleet operations. Autotrader's Co-Driver toolkit shows what governed, production-scale AI already looks like:

Our Co-Driver toolkit automates vehicle descriptions, image management, and retail workflows to help dealers advertise at scale.

Eleven thousand retailers have saved an estimated 200,000 hours collectively, with 96% accepting the first AI-generated description.

This proves that when AI is powered by quality best-in-class data, retailers can trust it to do the heavy lifting, removing repetitive tasks so they can spend more time focused on delivering an excellent customer experience.

That's what governed, production-ready AI looks like - not replacing people, but making them dramatically more effective.

Ian Plummer, CCO
Autotrader

// THE SECURITY ANGLE

enablis.

Why information security matters so much in fleet

Fleet isn't a typical target for cybercrime. In most industries, a data breach would stop at the network level, or the localised data level. In fleet businesses, cyber criminals can reach the vehicles themselves. A security problem can quickly become a safety and continuity problem, which is why information security discussions belong in the boardroom, not solely with the designated data security employee.

The 2025 Jaguar Land Rover attack forced a five-week shutdown and an estimated £1.9bn hit to the economy and KNP Logistics, a 158-year-old haulier, collapsed after a single weak password let ransomware in; around 700 jobs were lost. Ransomware on automotive and mobility firms doubled in 2025.

01

Cyber is now a safety risk

Documented attacks have used hijacked telematics to track and remotely disable vehicles, including cutting engines. For a fleet, a breach could mean downtime, data loss, and serious safety implications.

02

A new accountability gap

An agent can act wrongly with no attacker involved. Most fleets can't yet say which identity it acted under, what it was allowed to touch, or why it acted. Over half of organisations report no clear ownership of AI identities.

03

Telematics is personal data

Driver location, speed and behaviour fall under UK GDPR. The ICO is blunt: monitoring during permitted personal use is "seldom justifiable", and a data protection impact assessment is mandatory for high-risk monitoring.

04

The regulatory angle

UN R155 and R156 cyber rules came into effect for GB type approval for new vehicles from June 2026.

// INDUSTRY ANGLE - BVRLA

enablis.

BVRLA

How BVRLA is building AI capability across its membership

The trade body's own approach mirrors the discipline we argue for: start narrow, build capability and use progress to decide your next step:

Most fleet businesses are still at the early stages of AI adoption, often starting with simple productivity tools like drafting emails.

But that's changing fast - over the next year, we expect to see AI move into more operational areas, particularly around safety, from AI-supported fleet management to AI-enabled cameras.

At BVRLA, we're already building the foundations to support that shift. Our introductory AI course helps members build confidence in the basics - understanding capabilities and limitations, choosing the right tools, and using them responsibly, with proper attention to privacy, compliance and human ownership of outputs.

We're also using AI within our own learning platform to help shape and deliver content, and the next step is to make that far more personal - role-specific learning journeys that can even draw on audit findings, so that an amber or red result translates directly into the right learning and a way to track progress against it.

We're seeing real appetite from members for more advanced support too, particularly around compliance, and we're actively exploring how to extend what we offer to meet that demand.

Mark Pow, Head of Learning & Development
BVRLA

// BUILDING SECURE AI

How to decide what AI can do on its own, and what it can't

Doing this well has very little to do with running more pilots. It comes down to a deliberate, multi-step approach to data security and AI governance, with the guardrails built in from the very start rather than bolted on once something goes wrong. In practice, **five decisions determine whether an agent is safe to run in your fleet on its own:**

Put a human where the action is irreversible.

The test isn't "should a human review this?" It's "does a human reviewing this change the outcome?" Let an agent book a routine service slot on its own, but keep a person on the decisions that cost money or carry risk: authorising an expensive repair, charging a customer for damage, or taking a vehicle off the road.

Give every agent least privilege, a kill switch and an audit trail.

Scope exactly what each agent can touch: one triaging maintenance has no business in driver personal data or the payments system. Log why it acted, so you can stand behind the decision when a customer disputes a charge, an insurer asks questions, or the DVSA does, and be able to stop it instantly. The industry's own security standards now rank identity abuse and over-trusting a slick AI answer as the top agentic risks.

Start with one narrow, measurable, reversible workflow.

Maintenance work-order triage with human sign-off is the textbook first step: high volume, easy to measure, and simple to reverse. Prove the value there before you point an agent at anything customer-facing or safety-critical.

Fix the data foundation first before anything scales.

An agent is only ever as trustworthy as the data underneath it. In a fleet, that data sits across telematics, maintenance records, leasing and driver systems, so getting it governed and joined up is the precondition for everything else.

Earn each step up the autonomy ladder.

Move from an agent that suggests, to one that schedules, to one that acts across the fleet, unlocking each step only when the controls for the last are proven in production.

"We get our agents to check each other's homework. One writes the code, another independently validates it against a defined remit. If something's wrong, it's contained to that one agent, limiting the blast radius, instead of cascading through the whole team.

Lee Tucker, Lead Consultant
Enablis

// SECURELY MODERNISING LEGACY SYSTEMS

What it takes to be ready in the AI age

We can see the foundations matter immensely when working with AI. Once they're in place, the question turns to how you can modernise your operations without stalling the business, and how to sequence the work. Agentic AI raises the stakes here, because an agent inherits whatever it's given: the quality of your data, the gaps in it, and the controls around it. Get the foundations right and you can extend autonomy with confidence; leave them for later and every new use case carries the same weakness forward.

This is where we bring structure and direction. We help you move off platforms that can't scale, replace manual processes with governed pipelines, and put cataloguing, lineage and data-quality controls in place, so every agent works from data you can stand behind, and every step forward is one you can trust.

> **We're encountering some instances where fleets are starting to experiment with AI, which is very much to be encouraged in search of new efficiencies, but the data they are using is fundamentally flawed. It's always been critical that your fleet information is dependable but AI arguably makes it more important than ever.**

Peter Golding, CEO
FleetCheck

// CLIENT STORY

Identifying 50 AI opportunities for Nexus

enablis.

Spotlight:

enablis. x nexus

a fleet-sector story

Nexus, an industry-leading vehicle and plant hire business, wanted AI across the whole company, not just engineering. Over two months we ran workshops across seven teams (supply chain, finance, sales, legal, PMO, HGV and support), mapped the real workflows, scored the opportunities, and turned the backlog into working prototypes in a single 48-hour build. It started with where work slows, not a catalogue of tools, with clear ownership and sequencing, so nothing scaled without a governance line behind it.

7 teams assessed

50+ AI opportunities prioritised

5 prototypes in 48 hours

“What made the difference was having a team that could work across the whole business, not just with our engineers, but with finance, operations and contact centre. We went from conversation to deployed solution in a day. That’s the pace we needed.

Pete Ennis, CTO
Nexus

// OUR OWN BUILD

Building the Enablis operating platform

enablis.

We built our own operating platform with a team of AI agents in 26 days. We could build with pace and safety through embedding guardrails from day one, including a boundary checker agent and a full test suite that failed the build whenever it went out of scope. Our AI agents were then free to develop while our human-made controls ensured the build remained controlled.

26 days to build

380 improvements made by AI

2,606 AI-orchestrated tests

2 people steering the project

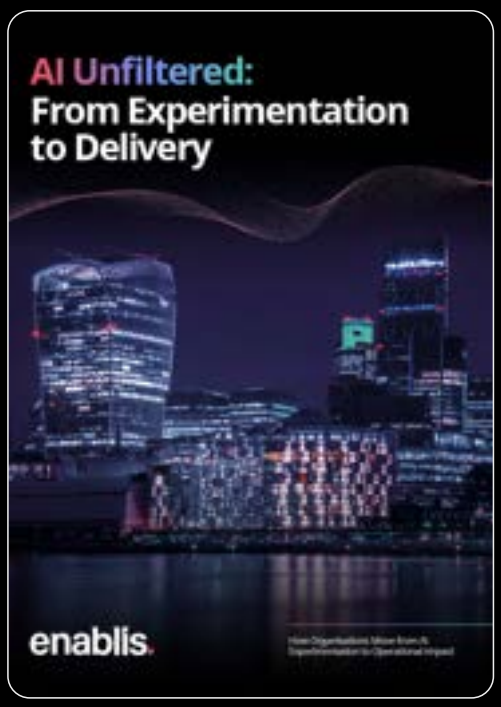

Read the full story in our **whitepaper**.

Governance isn't the most sexy word in the world... but for us governance is about the guardrails - support, tooling and reporting that allow a business to manage adoption effectively.

Ed Marshall, CTO
Enablis

// WHAT WE'RE BUILDING

enablis.

Tools we've built, ready for your organisation

We're not just advising on this. We've built the tooling to make governed AI faster and safer, and are rolling it into our own product range:

// FUSION

Our AI-native content and build platform. Collaborate with Fusion to decide how every site and product you build is structured, from a single site to a fully multi-tenant platform serving several clients at once. It runs the Enablis website today and client sites across multiple geographies.

// SKIPPER

Our AI-native operations engine. Skipper runs our timesheets, resourcing, pipeline and more - all in one place. It gives us one live picture of where every project and client relationship stands, from first conversation to invoice. An accelerator for any business.

// BOOK A 30-MIN DEMO

// HOW ENABLIS CAN HELP

We don't take over. We take your business further, faster.

We help fleet businesses turn AI ambition into governed, production-grade delivery, with the guardrails to run it safely. Wherever you're starting, there's a way in:

// ASSESS

We start with your business challenge, not the technology. Together we'll identify where AI can create the greatest value, explore the opportunity, and build a prototype where it adds confidence. You'll leave with a clear, evidence-based business case and a practical recommendation on what to do next.

// BUILD

We build production-ready AI solutions using proven accelerators, reusable blueprints and AI-native delivery. Everything is designed around your business, so you get the speed of repeatable delivery without compromising on what makes your organisation unique. Every engagement is focused on delivering measurable value from day one.

// TRANSFORM

Technology only creates value when people use it. We work alongside your teams to embed AI into everyday ways of working. Rather than creating dependency, we transfer knowledge as we deliver, building confidence and capability so your people can continue long after we've gone.

// SCALE

Once AI is proven, we help you deploy it securely across your organisation. From governance and integration to operating models and adoption, we make sure AI scales securely and delivers consistent value.

// BOOK A HALF-DAY WORKSHOP

hello@enablis.co
enablis.co
Leeds & London
LinkedIn: Enablis

enablis.

Profit from ideas.
The power of possibility.