Smart Freight Centre

June 2026

Agentic AI in Logistics: From Data Infrastructure to Operational Intelligence

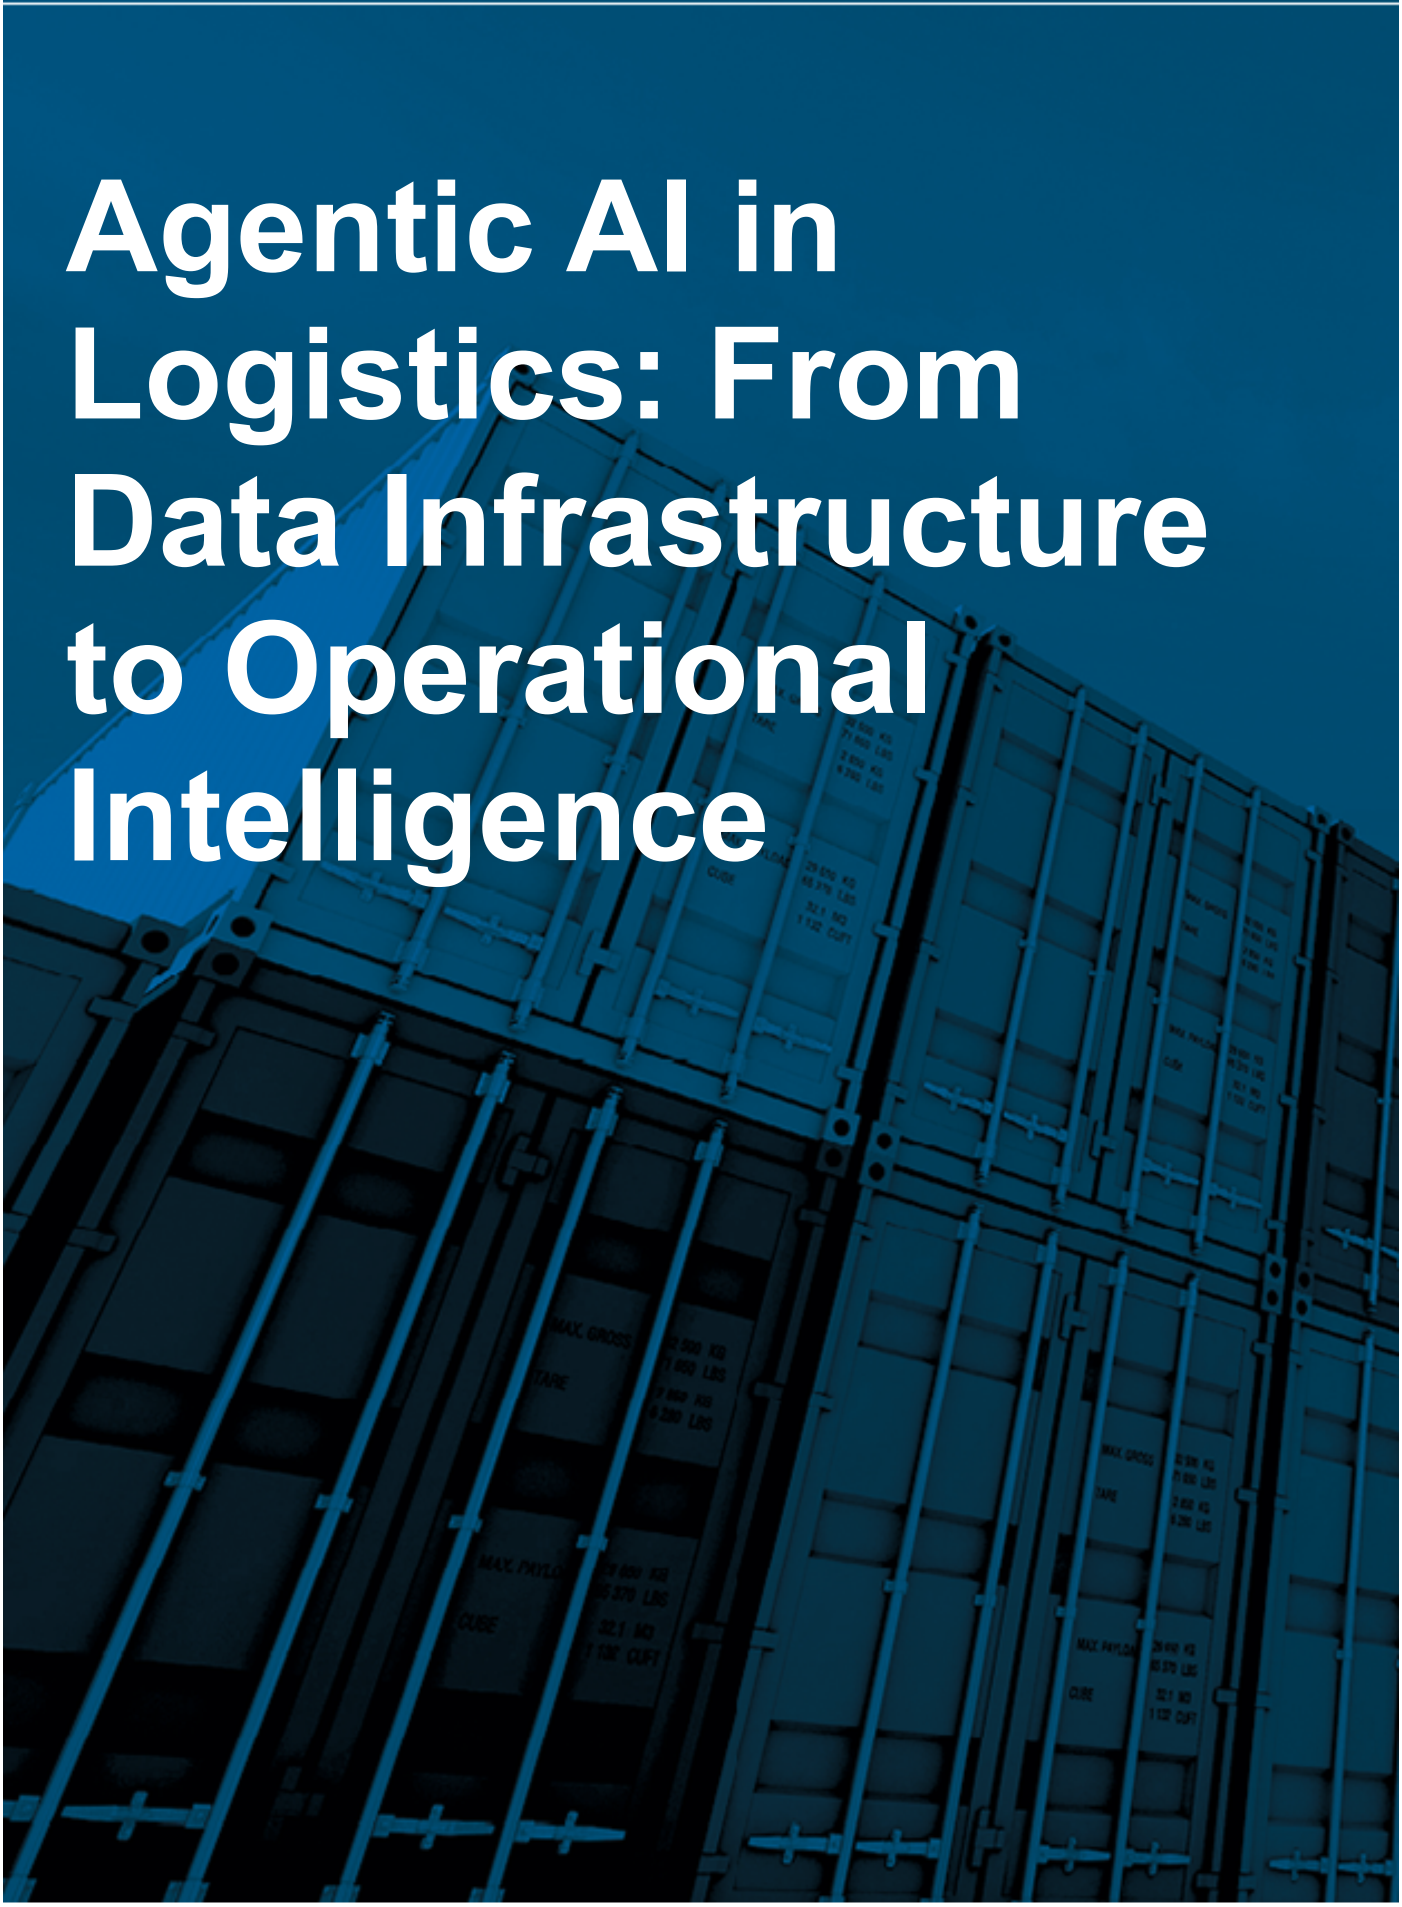

Acknowledgements

This whitepaper is co-authored by Smart Freight Centre (SFC), SINE, and two SFC Certified Tool Providers: Way Data Technologies and Pathfinder. It and expands on the session "AI Applications for Better Logistics Operations" which took place at Smart Freight Week 2026, at 11:30 CET on April 16, 2026. Smart Freight Centre moderated a panel featuring practitioners from all three contributing organizations. The details of the panellists and co-authors are:

- Violetta Matzoros, Senior Technical Manager at Smart Freight Centre
- Oscar Söderlund, Software Engineer at Way Data Technologies and Founding Investor at Greens Ventures
- Raimundo Henriques, Software Developer and iLEAP Project Lead at SINE Foundation
- Viktoria Kindesjö, Sustainability & Customer Success Lead at Pathfinder

If this is a topic to your organization's interest, Smart Freight Centre is prompting the reader to reach out to discuss ideas or suggestions for future work. Reach out to Violetta Matzoros at: violetta.matzoros@smartfreightcentre.org

Disclaimer

The information, views, and recommendations contained in this document reflect the collective expertise and practical experience of the contributors at the time of writing. They are intended for informational and discussion purposes only and do not constitute legal, regulatory, financial, or professional advice.
This document does not represent the official position of any individual contributing organization, nor does it create binding obligations on any party. The frameworks and standards referenced herein are subject to ongoing development and revision.

About Smart Freight Centre

Smart Freight Centre (SFC) is a globally active non-profit organization for climate action in the freight sector. Our goal is to mobilize the global logistics ecosystem, in particular our members and partners, in tracking and reducing its greenhouse gas emissions. We accelerate the reduction of logistics emissions to achieve a zero-emission global logistics sector by 2050 or earlier, consistent with 1.5° pathways.

About SINE Foundation

SINE Foundation is a non-profit that leverages digitization for the common good by developing open technology and standards. These enable companies to make sustainability-linked decisions and collaborate over highly sensitive data. SINE partners with others to establish global standards for sustainability data exchange. SINE is a co-sponsor and member of the iLEAP secretariat.

About Way Data Technologies

Way Data Technologies is an OEM-native fleet analytics platform built for service and delivery vehicle fleets. It connects directly to vehicle manufacturer APIs to bring primary data from mixed-brand fleets into one system, with no hardware required for data access. That verified foundation powers a decision layer above existing telematics and BI tools, helping fleets cut total cost of ownership, reduce downtime, and make smarter operational and electrification decisions. Way enables operators to move from dashboards to decisions, with AI agents that know the fleet, its vehicle health, recalls, and operational signals, so accurate data becomes a measurable result.

About Pathfinder

Pathfinder is an AI-powered logistics planning and decarbonisation platform developed by LOTS Group. Using real transport data, Pathfinder helps shippers, carriers, and logistics providers identify, simulate, and prioritise the most effective actions to reduce emissions, improve operational efficiency, and accelerate the transition to sustainable transport. By combining advanced analytics, artificial intelligence, and logistics expertise, Pathfinder enables organisations to make data-driven decisions that support electrification, alternative fuels, and smarter supply chain design.
www.pathfinder.lotsgroup.com

Contents

Executive Summary

Agentic AI is rapidly becoming a competitive differentiator in logistics. The size of the supply chain management software utilizing agentic AI capabilities is expected to grow over a 5-year horizon by more than 2500% to a total of $53 billion. Yet the gap between expectation and operational reality remains wide. This paper argues that a determining factor in whether an AI deployment succeeds or fails is not AI itself, but the underlying data: its availability, quality, and the processes and standards governing it.

Drawing on practitioner experience in fleet management, electrification planning, and data governance, this paper:

- Explains what agentic AI is and how it works, in simple terms;
- Identifies the most common misconceptions that lead to costly suboptimal deployments;
- Describes where agentic AI reliably creates value in logistics today;
- Outlines the governance structures organizations need before deploying AI at scale;
- Provides practical advice on AI deployment;
- Explains why open standards like, iLEAP, are necessary for AI to work reliably.

AI does not replace the need to access high-quality data, shared standards, and appropriate governance principles. It makes all three more urgent. Organizations that invest in these foundations now will be best positioned to capture the competitive advantages from agentic AI.

A Note on AI use and Sustainability

A legitimate tension exists between the rapid deployment of AI and the logistics industry's decarbonization commitments. AI infrastructure, in particular the data centers and specialized hardware required to run LLMs, consumes substantial energy. This is not a concern to dismiss.

The evidence cuts both ways. The hardware is improving quickly. Purpose-built inference chips now reach the market with substantially better performance per watt than the general-purpose processors of early deployments[1]. Open-source models run on standard hardware that could not be the case months ago. However, demand is outpacing supply, and the race to add capacity is pushing operators toward ever larger data centers on land and speculative designs at sea, each with its own land-use, water, and ecosystem effects[2]. Whether efficiency gains are absorbed by rising demand, and at what environmental price, is still an open question. Yet, several developments are worth noting. The specialized hardware now being designed for AI workloads is significantly more energy-efficient than the general-purpose graphics processors used in early large language model deployments.

The more immediate sustainability argument is one of proportionality; not every use case justifies the energy cost of running AI systems, and organizations should be deliberate about where AI genuinely delivers outcomes that outweigh its footprint. For the logistics industry specifically, the most defensible use of AI is in accelerating the transition to more efficient and cleaner transport. Using AI to identify fleet electrification opportunities, optimize route planning, or reduce empty running has a fundamentally different sustainability calculus than using it for more mundane tasks. For that use, the energy and carbon cost of the AI itself is likely to be a small fraction of the savings it enables.

Sources:

1. https://investors.broadcom.com/news-releases/news-release-details/openai-and-broadcom-unveil-llm-optimized-intelligence-processor
2. https://cleantechnica.com/2026/05/11/the-ocean-is-not-a-server-rack-panthalassa-peter-thiel-and-wave-powered-ai-compute/

1. Introduction

Logistics has always been a data-intensive industry. Logistics flows generate data about origin and destination, timing, payload, fuel consumption, route deviation, among others. Yet for decades, most of that data has remained siloed: carrier telematics systems that shippers cannot access, TMS platforms that fleet managers never see, energy consumption figures based on manufacturer assumptions rather than real-world performance.

Agentic AI does not dissolve this fragmentation. However, it does change what becomes possible when that fragmentation is overcome, and it makes the cost of leaving it in place far more visible. Gartner expects an exponential growth of supply chain management software with agentic AI capabilities from $2 billion in 2025 to $53 billion in 2030. This growth of more than 2500% over 5 years is mostly attributed to the increase in spending on AI agents[1].

Across the logistics sector, organizations are already taking action to integrate AI. A live audience poll conducted at Smart Freight Week (April 2026) found that most logistics professionals in attendance are actively working on AI integration[2]. The most active deployment areas reported were customer service, documentation, and data quality improvement processes.

This paper is written for logistics practitioners, in sustainability or operational roles, who seek to move beyond the hype and understand what agentic AI requires, what it reliably delivers today, and what must be in place for it to work at scale. The paper seeks to clarify common terms, address misconceptions, and inspire logistics teams to experiment with AI within their own organizations.

[1] https://www.gartner.com/en/newsroom/press-releases/2026-04-07-gartner-forecasts-supply-chain-management-software-with-agentic-ai-will-grow-to-53-billion-in-spend-by-2030

[2] The session *AI Applications for Better Logistics Operations* took place at Smart Freight Week 2026, at 11:30 CET on April 16, 2026. A poll conducted at the beginning of the session revealed that the vast majority of the audience is already integrating AI into their operations [see Appendix]. Given the title of the session, it is fair to assume that the audience was biased. Still, the results show a clear tendency.

2. Definitions

Before we dive into the definition of Agentic Artificial Intelligence (Agentic AI), we must first explain what a language model is, as it became the pervasive approach to implement Agentic AI.

A Language model is a specific machine learning approach in the domain of Artificial Intelligence. It is a statistics-based approach wherein a mathematical model "learns" the grammar, syntax, and semantics of one or more languages. Once such a model is trained, it can become capable of generating coherent-sounding and context-relevant outputs in one or more languages. Since 1990, language models have been developed using neural networks, but the results were modest.[3] The evolution of so-called "large language models" (LLMs) was made possible by the onset of large-scale pre-trained models, the development of transformers, and self-attention mechanisms.[4] LLMs are trained on vast amounts of existing data, and they excel at handling sequences of words and capturing patterns in text.

Agentic AI sounds complicated but is, at its core, built on three foundational concepts: a language model, system instructions, and tools. The language model enables the agent to process input and to generate new output. The system instructions give the language model its role, so that it behaves in a certain intentional way (see below). Last, but not least, tools (also see below) let the agent interact with its environment (see below).

Further reading: Google Cloud, "What is an AI Agent?" (2026); Anthropic, "Building Effective Agents" (2024).

System instructions are fixed, pre-defined text that the language model receives before any interaction. It defines the model's role, its domain knowledge, and how it should behave.

A fleet management agent, for example, understands vehicle data schemas, knows the regulatory environment for a given market, and responds in ways appropriate to that context. Different agents can be configured for different roles, each grounded in domain knowledge relevant to its function.

E.g.: System instructions for a fleet management agent
You are a fleet management assistant for a logistics operator in the European Union. You have expertise in EU road transport regulations, including driver hours rules under EC 561/2006 and tachograph requirements. You understand vehicle telematics data formats and can interpret odometer readings, fuel consumption logs, and GPS route data. You respond concisely and flag compliance risks when relevant. You do not speculate beyond the data provided.

Tools are where agentic AI becomes consequential. A tool is a capability you give an agent. For example, a specific action it can take in the world outside the conversation, such as booking a maintenance slot, sending an email, or running a route optimisation. Each tool is backed by software the agent does not run itself, such as an API, a routing engine, or a database query, that executes the actual task. An agent only has the tools you choose to give it.

In a logistics context, "tool" usually means a software product, such as an emission reporting tool, a transport planning tool, or a supply chain analytics platform. This term is used in a narrower, technical sense here. A tool is one specific action an agent can invoke. The two definitions connect

[3] https://www.edps.europa.eu/data-protection/technology-monitoring/techsonar/large-language-models-llm_en
[4] https://www.ibm.com/think/topics/large-language-models

directly though. When a purpose-built logistics tool is made “callable” by an agent, it becomes one of that agent's tools. Chapter 5. Why Standards are non-negotiable, returns to this, where agents work with specialized software by invoking it.

If you provide to a fleet management agent with the tool phrase "schedule_maintenance", it can output those words, and the system will call the pre-defined API and book a slot. Give it "run_route_optimization" and it can invoke a routing engine. This is what separates an agent from a chatbot. A chatbot produces text for a human to act on, an agent acts in the systems around it.

3. Challenges and Complications

Many misconceptions persist about what AI agents are and what they are capable of. Some misconceptions lead organizations to deploy AI in cases likely to fail; others cause them to hold back from deployments that could genuinely add value.

Often, unsuccessful deployments of AI agents are linked to a more fundamental issue: data asymmetries among and within organizations. Data asymmetries are in fact a common, yet fundamental challenge in the logistics sector. In this chapter, we therefore introduce the data asymmetry concept with a distinct logistics sector lens and then discuss common misconceptions in Agentic AI settings.

3.1 Data Asymmetries

To set the stage, the reader would benefit from understanding information asymmetry, as defined in economic theory. According to Akerlof (1970), information asymmetry takes place when one party in a transaction has access to more or better information than another. This leads to imbalances in decision-making and potential inefficiencies.[5] This is illustrated through the popular example of used car markets, where sellers possess more knowledge about car quality than buyers. This can result in lower-quality cars dominating the market. This dynamic shows how uneven information distribution can distort outcomes and undermine trust in exchanges.

In the digital era, data asymmetries occur in situations where one party has more or better data than another party, due to the fragmentation the industry operates on. In practice, it typically occurs among different IT systems *within* an organization, or *across* different companies and their IT systems. Data asymmetries are the result of a disparity in control of and access to data, especially when conflicting commercial interests hinder the flow of data between organizations. While data asymmetries are not new, the advent of AI makes them a far more serious challenge than before.

3.2 Data Asymmetries in the Logistics Sector

Data asymmetries in logistics are a crucial structural problem: shippers, logistics service providers, and carriers each hold data that the others would benefit from in data-driven decision-making scenarios. Yet, they are siloed behind organizational boundaries due to conflicting incentives that prevent it from being shared.

A very prominent example is the data asymmetry between shippers and the logistics value chain, affecting shipment level and fleet level data visibility and access. Shippers and logistics service providers typically have detailed visibility at shipment level, such as origins, destinations, volumes, weights, and timing constraints. What they rarely have is shipment execution-level data from the operating fleets: energy consumption, energy use, charging times, real dwell times, and route deviations. This data often exists in telematics platforms controlled by subcontracted carriers.

[5] Akerlof, G. A. (1970). *The market for "lemons": Quality uncertainty and the market mechanism*. *The Quarterly Journal of Economics, 84*(3), 488–500. https://doi.org/10.2307/1879431

Carriers and fleet operators face the mirror image. They have rich execution data from connected vehicles (fuel and energy consumption, exact charging durations, vehicle health, driver behavior) but often have limited shipment-level visibility; i.e. what their trucks are carrying, because that information is held in TMS platforms further up the chain.

In the context of electrification, a third data layer has become equally critical: charging infrastructure data, i.e., the location, power rating, reliability, and price of charge points. This information is currently scattered across dozens of operator apps and platforms, often unverified and inconsistent, particularly when it comes to whether a charger is certified for truck and trailer use or whether it reliably delivers its rated power across seasons.

Where does the impact of data asymmetries lie? The impact of these asymmetries is most visible in strategic planning. Fleet procurement, charging infrastructure investment, or network design are multi-year decisions informed by simulation and optimization tools. When real operational data is unavailable, these simulations fall back on industry averages, forcing conservative safety margins into every assumption. Over a planning horizon of five to ten years, those margins compound into significant over- or under-investment. Access to primary fleet data and charging data changes this: consumption can be modeled per truck and trailer combination rather than from manufacturer estimates and charging times can be grounded in operational measurements rather than rated specifications. Both vary across seasons, and accounting for that variation is the difference between right-sized investments and costly buffers. Logistics is a multi-party operation. With data sharing, following specific data policies to respect competitive advantages, actors in logistics can maximize the total benefits they can have and increase their overall competitiveness.

3.3 Data Asymmetries and Agentic AI

An agent with access to relevant, sufficient, and high-quality data can automate complex workflows, surface patterns across large datasets, or support high-quality decisions. An agent working on incomplete data will either fail visibly, declining to answer, or, more insidiously, hallucinate, potentially cascading into expensive mistakes.

It is also risky to let agents rely on data built on assumptions. If those assumptions are not explicitly documented, the agent may treat estimated values as facts and layer further inferences on top of them, leading to misleading recommendations. With the advancement of agentic AI deployment, the consequences of asymmetries are amplified. Agents could operationalize data gaps by producing recommendations or shaping decisions at speed and scale. Where humans might recognize uncertainty and pause, agents can propagate incomplete or biased inputs across workflows, turning isolated data gaps into systematic decision errors. The risk is therefore not just suboptimal decisions, but the speed and scale such decision making can reach, across operations.

3.5 Common Misconceptions

“AI will replace software"

The most common and costly misconception is the belief that generative AI will replace purpose-built logistics software: simulation engines, optimization tools, and electrification planning platforms. Consider high-stakes, multi-year decisions, such as fleet procurement, charging infrastructure investment, and network redesign. The right approach remains a structured methodology of simulation and optimization, with rigorous scenario analysis that tests different fleet sizes, charging configurations, seasonal variations, and financial assumptions etc. (Agentic) AI can accelerate parts of this process, helping to prepare data, structure inputs, compare scenarios, and communicate findings. However, it is not a substitute for the rigorous and purpose-built software that these decisions require.

The right model is not AI instead of specialized tools, but to leverage AI *with* specialized tools. An electrification planning platform, for example, can integrate AI agents that interact with purpose-built routing and energy modelling engines by invoking them, interpreting their outputs, and moving planners more efficiently through complex analysis. AI adds speed and accessibility; the underlying tools provide the analytical foundation.

"AI Can Fix Poor Data"

A closely related misconception is that AI can compensate for data quality problems; that once you introduce AI, the underlying data issues will resolve themselves.

AI can help structure incomplete inputs, flag obvious errors, and accelerate the preparation of messy datasets. These are genuinely useful capabilities. But AI cannot manufacture operational data accuracy or resolve underlying data completeness. An agent that confidently generates an inaccurate ETA could make an unfounded routing recommendation. This uncertainty could then be carried forward into its decision-making process. The principle of "garbage in, garbage out" still applies. With agentic AI deployments, the output is not just text, it is a decision, or the input to a decision. This makes it even more pressing to address data asymmetries and data quality issues for Agentic AI deployments.

"AI generates new knowledge"

Language models are the result of a statistical process. Whenever they produce outputs, they behave as a statistical machine. But language models do not "invent" or "discover" new knowledge; they *transform* the information at hand and what they were built upon. Consequently, if relevant data is missing in the model's training data or in its runtime context, the model will either acknowledge that gap or fill it with a confident-sounding fabrication.

This has two practical implications. First, any "new" data that a model produces will inherit the characteristics of the data it was trained on, meaning that AI cannot create trustworthy information in the absence of real underlying data. Second, the apparent novelty of AI-generated solutions is illusory. Models tend to be agreeable. Rather than challenging the user's assumptions, they will reflect and refine them. Thus, what might look innovative coming from the AI model might simply be the user's own idea in a new light.

This does not make AI unhelpful. Reframing information, surfacing patterns, and reorganizing existing knowledge can be highly valuable. But users should be clear about what the model is doing: it is transforming available information, not creating new facts or independent insight.

"AI Is Only for Software Engineers"

This misconception runs in the opposite direction, causing organizations to underestimate where AI can be deployed and who can benefit. As described in Chapter 2. Definitions, an agent's capability comes from its context: the data it has access to and the role it is given. This means anyone whose job involves making decisions based on data can work with an agent. A fleet manager reviewing maintenance priorities, a sustainability officer compiling emission reports, a procurement lead evaluating tenders. Each of these roles can be supported by an agent configured with the appropriate data and tools. Agentic AI shifts the relative relevance of roles: such that domain expertise and good data is now more important than years before when building new systems and applications.

"Agentic AI Makes Standards Obsolete"

If an agent can read and translate between almost any data format, why agree on a common schema at all? It might seem that, with agentic AI, standards became unnecessary overhead. This reveals a misconception regarding the role of standards. As more decisions are made by agents exchanging data with one another, a shared and unambiguous definition of each data point matters more, not less. An agent that infers case by case how to “translate” between formats adds potential points of failure that can propagate silently into a decision. This type of work might require the agent to, for example, convert units (including derived units) and decide whether two similar concepts are the same. Small mistakes in this phase of the data exchange can compound and lead to significant discrepancies. Standards such as iLEAP,[6] exist to remove ambiguity, which is not eliminated by agentic AI (see Chapter 5. Why Standards are non-negotiable). What changes with agentic AI is the cost of adoption. The barrier to standards was rarely the standard itself, but the work of making existing systems conform to it. Given the full technical specification, an agent can complete roughly 80% of the implementation autonomously in minutes, and the effort of integrating that work into existing systems drops alongside it. AI thus lowers the cost of adoption, not the value of the standard.

[6] The iLEAP project page is: https://ileap.global/

4. Practical Use Cases & Opportunities

Agentic AI is already deployed in production across logistics and beyond. In this chapter, some real-world examples are presented to draw inspiration from. The selection of use cases has been done with the “filter” of logistics sustainability and efficiency. As these use cases have a focus on optimal and efficient asset utilization and agentic AI deployment, they could enable the needed reduction in GHG emissions that the sector is aiming for with the net-zero 2050 target. In addition, some policy implications are presented, with the goal of raising awareness, and informing logistics practitioners.

4.1 Use Cases

Tender Analysis and Round-Trip Identification

Large numbers of logistics organizations still conduct tender analysis primarily in spreadsheets. AI can add immediate efficiency here: cleaning large datasets, visualizing flow data, and identifying overlapping lanes or potential round-trip opportunities far faster than manual analysis allows. When combined with purpose-built routing and optimization tools, the result is a faster and broader analysis. More scenarios are explored, more constraints are considered, and more options are presented to the decision-maker. Evaluating and choosing the scenario that satisfies operational and sustainability requirements becomes the main task the logistics practitioner is asked to complete, rather than spending most resources in data manipulation and visualization.

Fleet Management and Maintenance Planning

Fleet managers work in a context where the necessary data (vehicle state, maintenance history, energy consumption, and utilization patterns) is available in one place. This makes fleet management one of the currently most productive contexts for agentic AI. Consider a fleet manager responsible for 200 trucks across multiple depots. An agent can review telematics data overnight, flag vehicle anomalies, detect increased fuel consumption (finding the likely root cause), and draft maintenance recommendations with supporting evidence. It then presents a prioritized action list to the fleet manager, who reviews and approves it in the morning. The human expertise and judgment needed for this part of the process are the same as before, but no time is spent building the decision-supporting material. What previously took hours of manual analysis now takes minutes of focused decision-making. Maintenance is the most relatable example because every vehicle needs regular service and inspection, but the same pattern applies across fleet operations: cost tracking, utilization analysis, and compliance reporting.

Electrification Planning and Scenario Analysis

Planning the electrification of a fleet requires combining shipment flows, vehicle operational data, and charging infrastructure data into actionable investment recommendations. For this kind of planning, vehicle performance should be grounded in real operational data and adapted to the characteristics of the specific routes being assessed. When these inputs are available and can be connected in a structured way, AI agents can dramatically accelerate the analysis: preparing and structuring shipment data, screening transport networks for electrification-suitable lanes, comparing scenarios, and evaluating charging implications at scale.

What remains difficult is full end-to-end automation without human oversight. Where analysis depends on fragmented partner data, unverified charging information, or rapidly changing operating conditions, the uncertainty is too high for fully autonomous decisions. The most reliable operational model today is AI-accelerated human planning: agents handle the analytical heavy lifting, while the planner remains in the loop for final judgements on consequential decisions.

AI-powered development tools are enabling people without programming backgrounds to build their own internal tools and dashboards. A transport planner can build a custom reporting view, or an operations manager can create a scheduling tool, without waiting for IT resources. This accelerates internal innovation. But it also means more data flows and more decisions being made by systems that no one centrally vetted. Capturing this value at scale requires, however, robust governance and policies.

4.2 Policy Landscape

The European Union during the last few years has been developing its digital strategy to improve competitiveness in the industrial sector, logistics included. As part of the Digital Strategy, three key policies are important to know and monitor for companies that undergo digitalization transition simultaneously with their 2050 net-zero pathway. The aim of this chapter is to clarify the scopes of these Digital Acts and relate them, where relevant, to the environmental focused policies, like CSRD or CountEmissionsEU.

There are three digital strategy acts concerning data, data access, and AI applications:

1. Data Act.[7] It defines a framework for facilitated data sharing, where users of connected products have the right of access to operational product data in a standardized manner. This aspires to increase data-driven innovation and increase data availability. The scope of the act covers business-to-consumer, business-to-business, and business-to-authority data sharing. Data holders, companies that manufacture connected products, must set up data governance contracts with the users to define data sharing practices. In logistics this becomes particularly interesting as data sharing is required for the industry to assess operational and emissions performance. Lastly, it defines requirements on interoperability to facilitate data flow between sectors and EU Member States, following the Common EU Dataspaces principles. For the logistics industry specifically, this could be an interesting domain to follow, operationalized with the EMDS.[8]

2. Data Governance Act.[9] It introduces a framework to increase trust in voluntary data sharing for the benefits of businesses and citizens. Specific governance rules are proposed, and the role of data intermediates is key in its implementation. For logistics decarbonization, this act could become relevant in helping member states monitor the decarbonization pathways and societal health goals on regional and national levels by facilitating Business-2-Authority (B2A) data sharing.

3. Artificial Intelligence Act.[10] It introduces a risk assessment framework for case-based AI deployment. Logistics is not categorized as high risk per se, but instead, the risk of specific AI applications is evaluated. While most logistics use cases, such as route optimization or emissions modeling, are deemed low or limited risk, certain applications, particularly those affecting worker management or safety-critical operations, may be classified as high risk under the Act. This could become more relevant when scaling (agentic) AI deployment.

Relating to the data asymmetry described in Chapter 3. Challenges and Complications, a parallel can be drawn by two EU regulations which are pushing for transparency from opposite ends of the

[7] https://digital-strategy.ec.europa.eu/en/factpages/data-act-explained
[8] https://transport.ec.europa.eu/transport-themes/smart-mobility/creating-common-european-mobility-data-space_en
[9] https://digital-strategy.ec.europa.eu/en/policies/data-governance-act-explained
[10] https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai#1720699867912-0

logistics value chain. The CSRD raises expectations on shippers to report Scope 3 emissions with greater rigor and transparency,[11] which increases the importance of moving beyond defaults. The EU Data Act (Regulation (EU) 2023/2854)[12] pushes manufacturers to provide fleet owners with standardized access to their vehicle data. CSRD sets the reporting standard; the EU Data Act unlocks the data access to meet it. The EU Data Act is now in force, and yet the real opportunities it brings are not sufficiently widespread in the logistics industry.

Where GDPR provides individuals with the right to access personal data held by organizations, the EU Data Act grants individuals and companies the right to access data generated by connected products they own or use. For logistics operators acting as users of connected vehicles, the Data Act provides a right to access product data and related service data generated by those vehicles where such data are not directly accessible, requiring the data holder to make "readily available data … and the relevant metadata" accessible to the user without undue delay. Such data must be made available free of charge, "of the same quality as is available to the data holder," and in a "comprehensive, structured, commonly used and machine-readable format," and, where relevant and technically feasible, continuously and in real time (Article 4).

Fleet operators can also designate third-party platforms (analytics tools, AI systems) to receive this data directly from the OEM on their behalf.[13] Primary fleet data has historically been accessible only through OEM platforms, often with limited coverage of the variables fleet operators need. The EU Data Act requires that this data be made accessible to the fleet operator and transferable to third-party software providers via standardized technical interfaces. Access is provided under fair and reasonable conditions and is not fully free but typically limited to marginal cost for the user.
For AI deployments in fleet management and electrification planning, this is a significant shift. Many of the data asymmetries that have required safety margins and modelled assumptions in strategic planning can now be resolved with real primary data. The actual (measured) energy consumption of specific trucks, in specific climates, on specific routes becomes available as input to simulation models and agents.

As with GDPR, the market response to the EU Data Act varies significantly. Some OEMs have invested in commercial-grade APIs that provide clean, complete data at operational resolution. These are the manufacturers whose data is already powering fleet management and emission reporting applications today. Others have more ground to cover, with large volumes of not standardized data, and important operational variables still missing. For logistics operators, now is a good moment to start conversations with OEMs about data access under the EU Data Act. The primary fleet data available under the regulation can be a valuable strategic asset, and great opportunities can follow from early engagement.[14]

[11] CSRD regulation: https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32022L2464
[12] Regulation (EU) 2023/2854 of the European Parliament and of the Council of 13 December 2023 on harmonized rules on fair access to and use of data. OJ L, 22.12.2023.
[13] EU Data Act, Article 5(1): "upon request of the user, data holders shall make data accessible to a third party designated by the user."
[14] For a practitioner perspective on the technical challenges of accessing and using primary vehicle data under the EU Data Act, see: Söderlund, O. (2026). "From CAN Bus to Carbon Report: The Primary Data Challenge in Logistics." LogR, Nr. 01/2026, pp. 14-15.

5. Why Standards are non-negotiable

A recurring theme across every discussion of AI deployment in logistics is the role of standards. We argue that there is a misconception that AI would make standards less necessary by automating data translation. Instead, standardization of semantics and data access enables effective and more trustworthy agentic AI use cases at a lower cost and at a greater scale.

5.1 AI Amplifies Good Standards

When data flows between systems in a common standard, AI agents can process it efficiently, accurately, and at scale. When data requires "translation" e.g., unit conversions, schema mapping, or terminology reconciliation, AI can help, but each translation step introduces risk. Unit conversion errors have caused significant divergence in emissions calculations and other logistics metrics. An agent that makes a unit conversion error confidently, at scale, across thousands of shipments, is more dangerous than a human making the same mistake in a single report.

Standards eliminate these "translation" steps. They allow agentic AI to focus on higher-value tasks rather than data preparation and format harmonization. Without standards, agents have more work to do before they can begin to be useful, and more opportunities to make mistakes along the way.

5.2 iLEAP: Open Standard as Enabling Infrastructure

Since 2022, Smart Freight Centre and SINE Foundation have been working with logistics actors on iLEAP: the open, community-driven standard for the peer-to-peer exchange of logistics emissions data. The iLEAP Technical Specifications[15] define everything that is needed for AI agents to make use of emissions data: a data model that is aligned with the GLEC Framework[16] and ISO14083, plus an HTTP REST API so that agents can access the data seamlessly from every implementing software solution.

When carriers, shippers, LSPs, and freight forwarders structure their logistics emissions data according to iLEAP, the data gaps between them narrow. Agents given access to iLEAP data can operate with confidence that they are working with standardized information, rather than spending resources reconciling inconsistent formats or units.

Implementing iLEAP is itself an AI use case. An agent equipped with the full iLEAP Technical Specifications can autonomously implement a substantial portion of the requirements, drastically reducing time and technical overhead. AI and standards are mutually reinforcing. Standards ensure that AI is provided with meaningful, well-defined tasks to work on, while ensuring data quality and avoiding risky ad-hoc "translation."

iLEAP is enabling infrastructure for AI-powered logistics. It is not a parallel workstream, nor something AI can displace. It is a prerequisite that can leverage AI to accelerate adoption.

15 https://specs.ileap.global/
16 GLEC is the globally recognized methodology for logistics emission accounting and reporting, it serves as the primary guidance for companies who wish to implement and comply with ISO 14083 standard.

6. Recommendations for Logistics Organizations

This chapter provides some practical tips for organizations that are planning or are already advancing their roadmap to agentic AI integration.

Before You Deploy

Map your data gaps. Identify the data you have, what you need, and what the gap between them means for the AI use cases you are considering. Do not deploy agents in contexts where the gap will force them to operate on assumptions. The compounding of unverified assumptions is one of the most common failure modes in logistics AI.

> Designing an agentic system means deciding which tools to give agents, which actions require approval, and under what conditions. The more consequential the action, the more important it is that the agent is working with complete, high-quality data.

Establish data governance. Approach it as a structured framework anchored in three principles: data quality, data sovereignty, and accountability. Organizations can start by defining and documenting data provenance, collection methods, validation cycles, and known limitations for every dataset used by AI. Governance shall also reflect ownership and usage rights. Explicit agreements are required on how shared data can be used, including for AI-driven decision-making, especially across organizational boundaries and jurisdictions. Finally, assigning accountability is key. This includes responsibility for input data quality, system design and configuration, and the consequences of AI-supported decisions. As agents increasingly interact with other systems, governance should anticipate security risks, ensuring auditability, traceability, and clear liability structures.
Update your data sharing agreements. Existing agreements may not cover AI use of shared data. Review and update consent and sharing frameworks before deploying agents that process partner data.
Assess your EU Data Act entitlements. If you operate connected vehicles in the EU, engage with your OEM about data access. The primary fleet data you are legally entitled to is likely more valuable to your AI programs than any synthetic alternative, and it is available now.

When You Deploy

Start with local data contexts. Deploy AI first in use cases where all relevant data is in one place and under your control. Build confidence in the technology and your organization’s ability to work with it before extending to cross-organizational settings.
Keep humans in the loop for high-stakes decisions. Agentic AI is most reliable today as an accelerator for human decision-making, not a replacement for it. Design workflows where agents handle the analytical heavy lifting and humans retain final judgement on consequential decisions.
Adopt open standards. Implement iLEAP and other relevant (open) data standards. This reduces integration friction, enables data to flow across the chain, and makes AI deployment faster, cheaper, and more reliable at every step.
Consider privacy-preserving technologies (PETs) for cross-organizational data collaboration. PETs enable cross-organizational data collaboration by preventing sensitive information from being disclosed. This unlocks use cases that would otherwise be commercially impractical. Multi-Party Computation (MPC), for instance, lets multiple organizations run computations on their combined data (e.g., running benchmarks) without exposing it. This

technology is available today: SINE's MPC engine Polytune is already deployed by German municipalities and runs directly in the browser.[17]

As You Scale

Build AI literacy across your organization, not just in your IT team. The organizations capturing the most value from AI are those where people across functions understand how to work with agents: how to give them the right context, how to evaluate their outputs critically, and how to recognize when to escalate to human judgement. This is a capability that can be developed incrementally, starting with a single tool, a single delegation, and a single use case.

Contribute to the data commons. Data asymmetry in logistics is a collective problem that no single organization can solve unilaterally. Participate in standardization efforts, data-sharing consortia, and collaborative infrastructure that makes trustworthy data available across the industry. The organizations that contribute to this infrastructure benefit from it directly and the industry as a whole moves faster. Diving deeper into the Data Governance Act (see Chapter 4.2 Policy Landscape) could lead to a use case where your organization could "give back to society" by sharing data that relate to Sustainable Development Goals or other data transparency initiatives relevant to your operating region.

[17] See https://sine.foundation/library/mpc-vaccination-check and try Polytune for yourself at https://benchmarking.sine.dev.

7. Conclusion

Agentic AI is not a future technology. It is available today, it is being deployed, and it is creating measurable competitive advantages for the organizations that understand how to use it well. The key takeaway of this paper is straightforward:

AI is as good as the data it has access to, the governance that surrounds it, and the standards that enable data access for agentic AI

Organizations that invest in data infrastructure, establish clear governance frameworks, and adopt open standards are building the operational foundations that will determine their competitiveness for the next decade, where AI is included. Industry associations, standard-setting organizations, and NGOs like the Smart Freight Centre and SINE Foundation play a decisive role in this shift. No single organization can solve interoperability or data asymmetries alone. The frameworks and standards that enable data flow across the logistics system are collective achievements, requiring sustained collaboration to develop, maintain, and adopt at scale. Organizations that engage with these initiatives, contribute to their development, and implement their outputs are helping to build the infrastructure that the entire industry depends on.

This does not mean organizations must halt their agentic AI deployments and wait until every data challenge is resolved before acting. Valuable AI use cases already exist today in areas such as data preparation, workflow automation, scenario analysis, reporting, and decision support. The organizations that move fastest will often be those that start with bounded, high-value applications in existing workflows, while strengthening the data, governance, and standards foundation needed for broader deployment over time.

The trajectory for agentic AI in logistics will be increasingly shaped by how effectively organizations engage with the emerging data and policy landscape. As outlined in Chapter 4.2 Policy Landscape, the combination of the EU Data Act, Data Governance Act, and CSRD is asking for higher access to primary operational data and strengthening transparency requirements. This creates a pathway for more reliable, data-driven AI applications, particularly in areas such as fleet management and electrification planning. The extent to which this opportunity materializes will depend on how proactively organizations secure access to primary data, align their data governance with these frameworks, and integrate these inputs into operational decision-making processes.

This paper has focused on the data and operational conditions required for effective AI deployment, but several areas remain outside its scope. In particular, the broader implications of the EU digital strategy acts, including the Data Governance Act and AI Act, merit deeper analysis as their implementation evolves. Further work is also needed on how organizations structure governance for digitalization pathways. This includes which teams own AI processes, training, evaluation, validation, and how accountability translates to AI-supported decisions. Finally, the role of privacy-enhancing technologies in logistics remains underexplored, particularly in enabling secure cross-organizational data collaboration and unlocking new forms of data sharing that support sector-wide decarbonization. These topics could be addressed in future publications to further support the sector's transition.

Appendix

The results of the poll from the session “AI Applications for Better Logistics Operations” took place at Smart Freight Week 2026, between at 11:30 and 12:30 CET on April 16, 2026. A Mentimeter poll conducted at the beginning of the session revealed that the vast majority of SFC members and partners is already integrating AI into their operations (Figure 1). Given the title of the session, it is fair to assume that the audience was biased. Still, the results show a clear tendency that is confirming the industry trends. In addition, the most popular applications and processes are presented in the form of a word cloud (Figure 2).

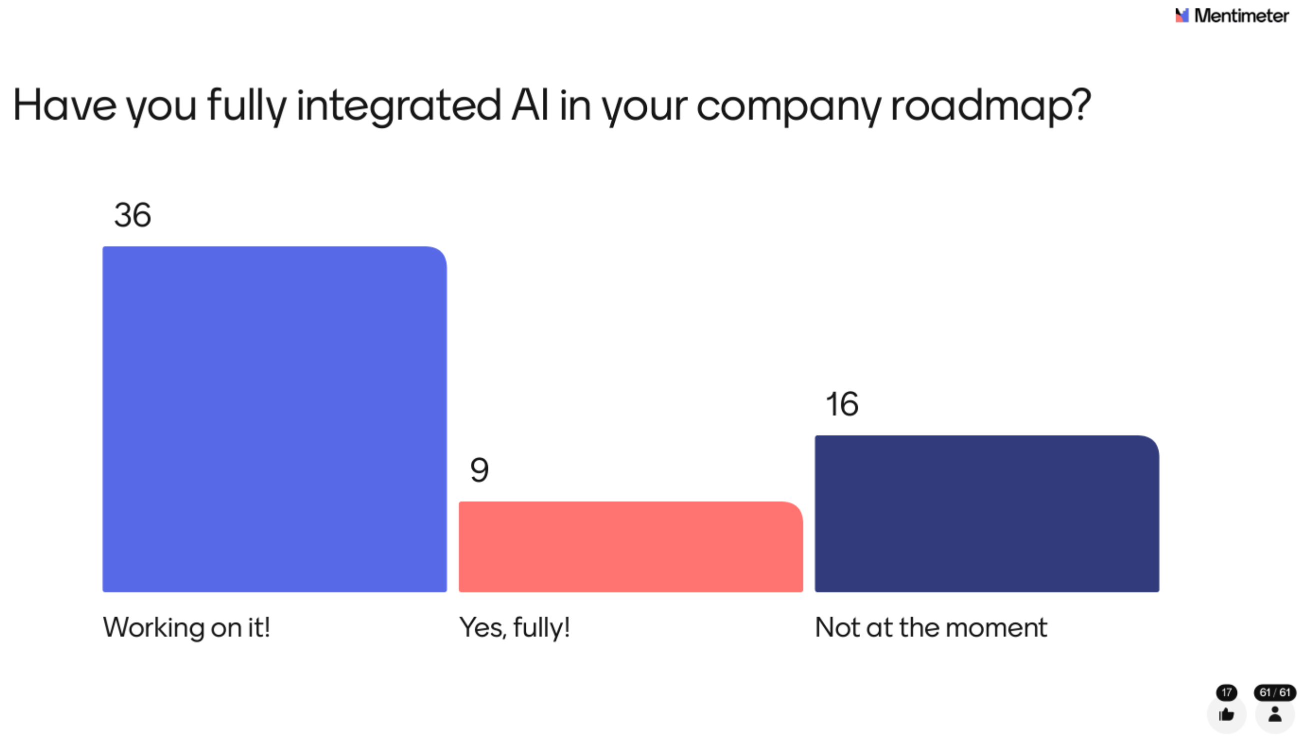

Figure 1 : Extent of AI integration in member’s strategic roadmap.

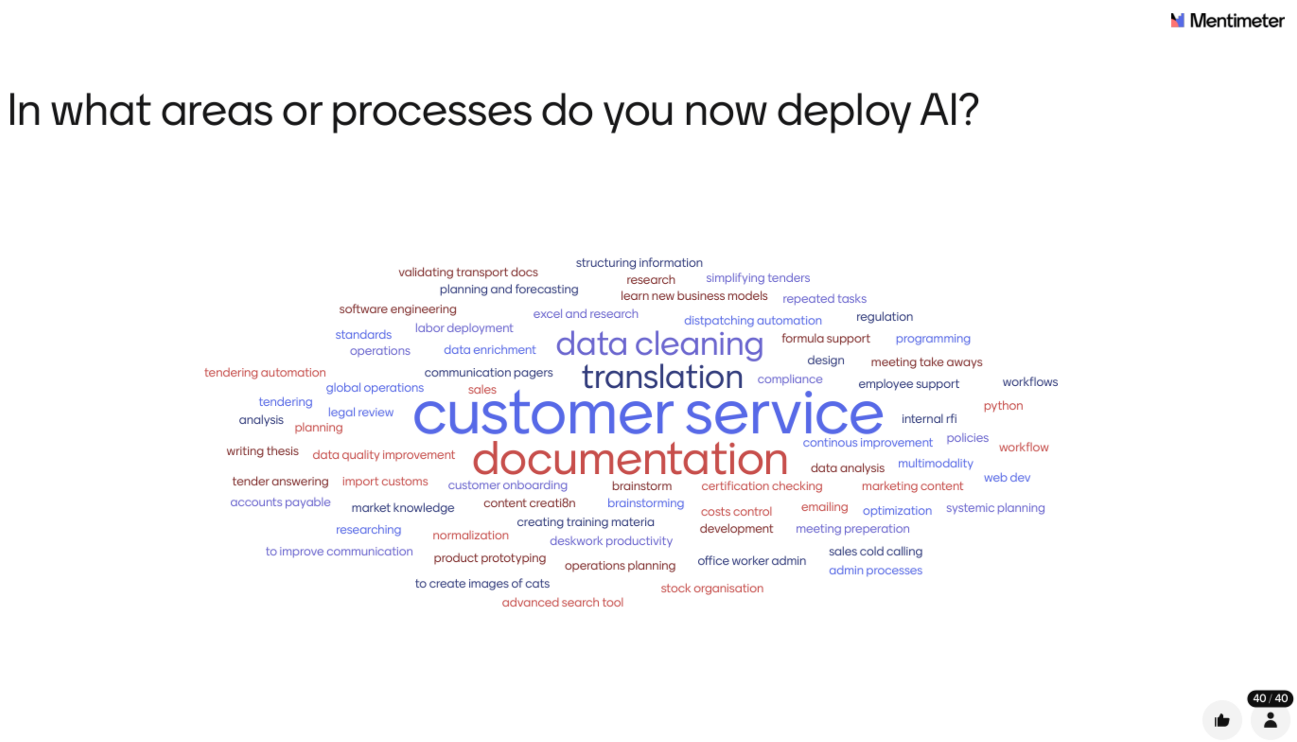

Figure 2: Most popular processes where AI is currently deployed.

How can you accelerate your decarbonization journey?

June 2026

Suggested citation: [Smart Freight Centre, SINE Foundation, WAY Technologies, Pathway] [Agentic AI in Logistics: From Data Infrastructure to Operational Intelligence] 2026

Acknowledgements
This whitepaper was written by Violetta Matzoros of Smart Freight Centre and co-authors Oscar Söderlund, Software Engineer at Way Data Technologies and Founding Investor at Greens Ventures, Raimundo Henriques, Software Developer and iLEAP Project Lead at SINE Foundation, and Viktoria Kindesjö, Sustainability & Customer Success Lead at Pathfinder.

Smart Freight Centre (SFC) is a globally active non-profit organization for climate action in the freight sector. Our goal is to mobilize the global logistics ecosystem, in particular our members and partners, in tracking and reducing its greenhouse gas emissions. We accelerate the reduction of logistics emissions to achieve a zero-emission global logistics sector by 2050 or earlier, consistent with 1.5° pathways.

Contact
Smart Freight Centre
Keizersgracht 572, 1017 EM, Amsterdam, Netherlands
www.smartfreightcentre.org
info@smartfreightcentre.org