WHITEPAPER

The Practitioner's Guide to AI in Logistics

WHY AI MAKES DISCIPLINED TECHNOLOGY DECISIONS MORE IMPORTANT

JBF
design to deliver

Table of Contents

Introduction

AI is changing logistics technology quickly, but the decisions surrounding its investment and adoption are like those required for any significant technology initiative. The business still needs to define what it is trying to improve, determine whether the solution is appropriate for the problem, and establish how success will be measured and managed over time. AI can amplify the consequences of weak assumptions or unclear ownership because its outputs can increasingly influence operational decisions without the same level of direct human review. This makes established disciplines around business case development, ownership, governance, and measurement increasingly important.

That need for discipline is already visible in how AI investments are being pursued. 78% of organizations are pursuing AI without a documented plan, while only 3% return afterward to compare results with the business case that justified the investment.[1] This is a familiar challenge in technology investment: interest in a capability can move faster than the practices used to evaluate whether it fits the business and to verify whether it delivered what was expected.

AI does not require shippers to abandon the decision processes they already use for significant technology investments and replace them with an entirely new playbook. It requires greater attention to detail in applying those disciplines because the technology introduces different ways for assumptions to carry into the operation. A poorly defined use case can direct investment toward a problem that was never material, unreliable data can shape outputs that appear credible, and an unclear operating model can leave responsibility unresolved when the system begins influencing decisions that were previously made by a person.

This guide follows that decision process from evaluation through ongoing operation. It examines how to understand what a vendor is offering, determine which use cases are worth pursuing, build a business case that reflects the way AI costs and benefits behave, define appropriate limits on autonomy, and continue testing the solution against the business objectives that justified the investment. The technology changes across those stages, but the need for sound logistics, business, and technology judgment does not.

[1] Source: JBF Consulting's 2026 *AI Without a Plan Survey Report*

1. Finding and Prioritizing the Use Cases Worth Pursuing

Technology-first thinking creates one of the most common sources of wasted effort. A technically successful pilot can still produce little business value if the use case was selected because AI could perform the task rather than because the task represented a meaningful problem.

One example from JBF's advisory work involved a team building and training an elaborate AI module to generate and summarize a complex report. They cheered their success, but the report never got used for making faster strategic decisions, as intended, because they never engaged the correct stakeholders. In turn, those stakeholders missed the window to provide key requirements and generate training for the end users. The AI worked, but the problem it solved did not justify a new investment.

85% of surveyed supply chain leaders reported that AI opportunities are selected without a consistent, formal process, and 67% identified building or validating the business case as a primary barrier to progress.[1] A useful starting point is to determine whether solving the problem would improve an outcome the business already cares about and whether the current process, data, systems, and ownership can support the proposed change. If the value is material but the foundation is not ready, the next investment may be process standardization, data remediation, or traditional automation rather than AI.

The process view is particularly important in logistics because the work that appears straightforward at a high level often contains the exceptions that determine whether automation will hold up in practice. A transportation process may follow one standard path for most customers while requiring different handling for a specific customer, site, mode, or regulatory condition. In other cases, there may be no consistent rule at all because teams have developed different ways of handling the same issue over time. Mapping the work as it is performed makes those conditions visible before they are embedded into a model or handed to an agent as part of an objective.

[1] Source: JBF Consulting's 2026 *AI Without a Plan Survey Report*

Use-case value also extends beyond eliminating manual tasks. Labor savings are usually easier to quantify, but many of the more meaningful logistics opportunities involve improving the quality or speed of a decision rather than removing the person who makes it. AI may help a planner identify an exception earlier, evaluate a larger set of options, or interpret information that would otherwise take too long to review manually. The business case still needs to define what changes because of that improvement, but the absence of a direct headcount reduction does not mean the use case lacks value.

The discipline in use-case selection is therefore not about finding as many opportunities for AI as possible. It is about understanding the operation well enough to distinguish where AI adds something the existing process cannot do efficiently from where a simpler change would solve the problem with less cost and complexity.

2. Reading an AI Vendor Demo

Evaluating an AI capability starts with understanding what is being offered and what will be required to support it after implementation. That has become more difficult as AI has developed into an umbrella term for capabilities that range from long-established automation to machine learning, generative AI, and agentic AI. The terminology matters less than the underlying behavior of the technology, because different capabilities create different expectations for the data they require, the way they are maintained, and the degree of oversight the shipper will need to provide.

JBF was recently briefed on a vendor's new AI feature that, after further questioning, turned out to be a predefined response to a predefined condition. The capability may still have been useful, and the labeling did not necessarily reflect an intent to mislead. The practical issue was that evaluating the feature as AI could cause a buyer to make assumptions about what the technology learns, how its behavior may change, and what ongoing responsibilities come with it that did not match how the feature worked.

The reverse can create a larger operating risk. Machine learning uses historical patterns to generate predictions, while generative AI interprets information and produces new content. Agentic AI goes further by pursuing an objective across multiple steps and, depending on its permissions, taking action. A prediction that still requires a planner to decide what to do is operationally different from an agent that can change an appointment, tender a load, or communicate with a customer without waiting for approval, even when both capabilities appear under the same AI heading in a sales presentation. Our 2026 survey illustrates how broad the terminology has become.

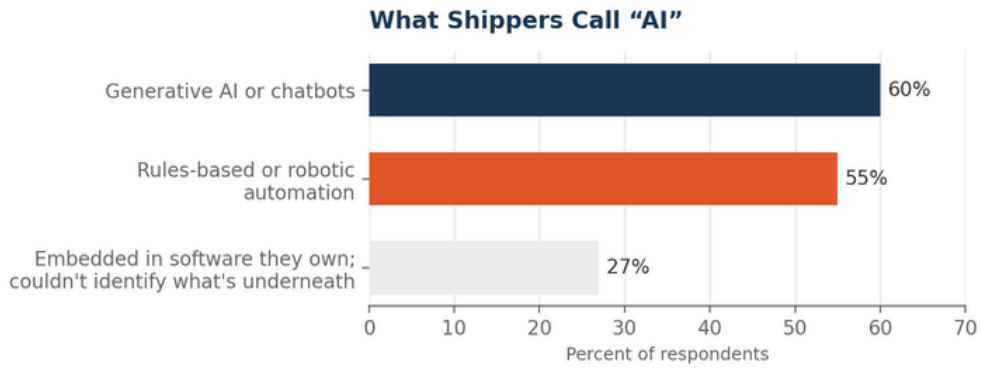

Source: JBF Consulting's 2026 *AI Without a Plan Survey Report*.
Categories are not mutually exclusive; respondents could name more than one.

Those responses do not suggest that logistics leaders need to become data scientists before evaluating a vendor. They do suggest that the buying process must move beyond the label and establish how the capability works in the specific operating context in which it will be used.

A strong evaluation therefore needs to connect the demonstration to what happens after the demonstration. The buyer should understand what information the capability relies on, how exceptions are handled, what changes when the underlying conditions change, where human approval remains necessary, and which responsibilities stay with the vendor versus transfer to the shipper. These are familiar technology evaluation questions, but AI raises their importance because the effect of an incorrect assumption may not remain contained within a configured workflow. The more the technology can infer or act, the more important it becomes to understand the business judgment that has been translated into the system and how that judgment will be monitored over time.

3. Modeling AI ROI When Costs Move With Usage

Selecting the right use case does not resolve the investment decision unless the business can define what success should look like and how the economics will be tested.

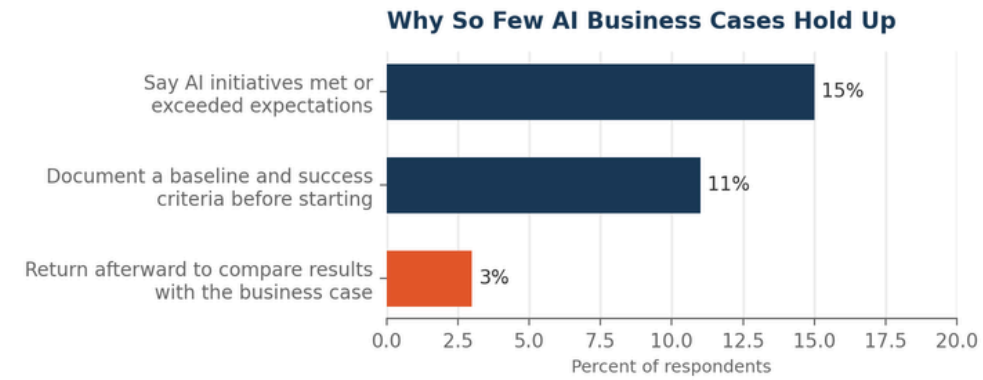

Source: JBF Consulting's 2026 *AI Without a Plan Survey Report.*

Without that baseline and follow-through, an organization may know that a capability was deployed without being able to say whether the investment produced the value that justified it.

The underlying discipline is the same one required for any technology business case, although AI introduces cost behavior that can make weak assumptions more visible as usage grows. Many logistics software agreements use relatively predictable subscription structures, sometimes with pricing tied to annual freight under management, shipment volume, or another negotiated tier. AI services may instead introduce charges based on tokens, application programming interface calls, compute usage, or other consumption measures that move with the volume and complexity of the work being performed. A pilot that looks inexpensive at limited usage can therefore produce a different cost profile once the capability is embedded into daily operations.

The benefit side of the case requires the same care. When AI removes a discrete manual task, labor hours can be estimated reasonably well, but hours saved do not automatically become financial savings. If headcount, overtime, or outsourced labor do not change, the organization has created capacity rather than reduced cost, and the business case should identify how that capacity will be used and how the benefit will be measured.

A published Coca-Cola North America example, in which an AI agent reduced the response time for 'where is my order' inquiries from a 90-minute service commitment to seconds, illustrates the point. The value can include labor capacity, but faster response may also affect service, escalation volume, and the customer experience, which require different measures than an hourly labor rate.

AI pricing also adds uncertainty to longer-term projections because usage patterns and commercial models are still evolving. The business does not need perfect visibility into future token or compute pricing to make an investment decision, but it should understand how sensitive the case is to higher usage, different transaction patterns, or a future rate change. Treating those factors as variables is more defensible than presenting a pilot's current unit cost as though it will remain fixed through a multi-year operating period.

A credible AI business case therefore starts with a documented baseline, states what should change and by how much, models the cost in a way that reflects how the tool will be consumed, and establishes when the result will be reviewed. None of those requirements is unique to AI, but the combination of variable cost, less predictable operating behavior, and faster scaling makes it easier for a weak assumption to compound before the business recognizes that the economics have changed.

4. Where Autonomous AI Stands Today

The discussion becomes more consequential as AI moves from generating information to acting. Much of the current attention on agentic AI centers on how much work an agent can perform autonomously, but the more important operating question is how clearly the organization has defined what the system is trying to accomplish and the conditions under which it should stop and wait for a person.

Logistics has dealt with a version of this issue before. Early route optimization technology could produce plans that were mathematically stronger than a dispatcher could build manually, yet many organizations struggled to realize the expected value because the operating judgment behind the plan had not been translated completely into the configuration. Service commitments, customer preferences, equipment constraints, cost tradeoffs, and exceptions all had to be represented accurately before the business could trust the output without checking every result. Agentic AI increases the importance of that translation because the system may adapt how it pursues an objective as conditions change rather than simply executing a configuration until someone changes it.

This changes the role of the person around the process. Work that was previously performed by hand can shift toward defining the correct outcome, establishing the boundaries within which the system can act, watching whether its behavior continues to reflect those expectations, and intervening when it does not. That responsibility is easy to overlook if governance is treated primarily as access control, security, or a policy document. The operational owner also needs a way to recognize when the system is technically functioning but producing a result that is wrong for the business.

The market itself remains earlier in this transition than the autonomy narrative can suggest. 23% of supply chain leaders reported having AI in live production somewhere in the organization, while 40% remained in prototype or pilot.[1] Public examples from companies investing heavily in AI also tend to show agents scoped to defined functions rather than an entire logistics operation being handed over without human involvement. C.H. Robinson, for example, has described agents supporting activities across the quote-to-cash process while continuing to position human insight and AI as complementary parts of the operating model.

[1] Source: JBF Consulting's 2026 *AI Without a Plan Survey Report*

The governance data shows where the next operating challenge sits. Only 10.5% of survey respondents reported a formal, actively maintained governance framework after AI goes live, while nearly a third were unsure whether one existed. A workable governance model needs to establish who owns the system's behavior, how performance will be reviewed against the intended objective, which exceptions require human judgment, and what changes to the process or environment should trigger a reassessment of the original design. The purpose is not to prevent the system from adapting; it is to make sure the organization can see whether the way it is adapting remains consistent with the outcome the business approved.

5. Sustaining the Result After Go-Live

Go-live does not end the investment decision because the conditions that made an AI use case appropriate can change. Volume changes, customer requirements evolve, data sources shift, models are updated, and teams redesign the processes around the capability as they learn where it performs well and where it does not. A business case that was reasonable at implementation can therefore weaken over time even when the technology continues to operate as designed.

Traditional logistics systems also require maintenance, but AI can introduce additional ways for performance to drift from the original expectation. A prediction can become less accurate as the patterns in the underlying data change, a generative model can behave differently after an update, and an agent can continue pursuing the stated objective while the business conditions around that objective have changed. Monitoring should therefore extend beyond system uptime and adoption to whether the outputs and actions still support the business result that justified the investment.

This is where domain expertise becomes particularly important. The technology can measure many aspects of its own performance, but it cannot independently determine whether the business is still asking it to solve the right problem. Someone who understands the logistics operation needs to recognize when an apparently correct result conflicts with a service commitment, customer requirement, operating constraint, or policy that the system has not represented adequately. That judgment may sit inside the shipper, with a technology provider, with an outside advisor, or across several of those groups, but it needs to be assigned rather than assumed.

The same principle applies when an organization begins to scale AI. A successful use case should not automatically become the template for every adjacent process because the conditions that made the first implementation appropriate may not exist elsewhere. Scaling should repeat the same evaluation at the next level, confirming the problem, the value, the readiness of the process and data, the ownership model, and the measures used to determine whether the capability is working. The discipline does not become less necessary because the organization has already proven that AI can work in one part of the operation.

Sustaining AI therefore looks less like maintaining a finished software implementation and more like maintaining the alignment between the business objective, the operating process, and the technology that is acting against it. The organizations that can make that alignment visible are better positioned to expand the use of AI without losing track of why each investment was made in the first place.

Conclusion

AI is introducing capabilities that logistics organizations have not had before, but it does not remove the need for the disciplines that have always determined whether a technology investment creates value. The business still needs to understand the problem, evaluate the solution against the way the operation works, define the economics, establish ownership, and verify the result after deployment. AI makes those decisions more consequential as the technology moves from assisting a person with information to influencing or acting within the operation.

The practical implication is not that every shipper needs to build a separate AI organization or develop deep internal expertise in every model and architecture. The organization does need access to enough logistics, technology, and implementation judgment to challenge assumptions before they are embedded into the solution and to recognize when the result no longer reflects what the business intended. That capability can come from internal teams, external partners, or a combination of both, but it cannot be replaced by the technology itself.

JBF supports shippers at those decision points without assuming that AI is the answer before the problem has been understood. Our strategy work helps define and prioritize use cases, assess the process and data conditions that will determine whether they are viable, and build the business case before an investment is made. Our vendor advisory work provides an independent view of what a solution does and how well it fits the requirements. When a capability moves forward, our delivery teams help translate the operating requirements into the process, integration, governance, and adoption work needed to put it into use, while continuous improvement keeps the original business case connected to what the solution is producing after go-live.

That approach is not fundamentally different from the way JBF evaluates other significant logistics technology decisions. AI changes the capability, the speed, and in some cases the level of autonomy, but the investment still depends on the quality of the business decisions surrounding it.

About the Authors

Tara Buchler is Principal, Strategy at JBF Consulting, bringing more than 20 years of experience at the intersection of logistics operations and enterprise supply chain software. She partners with shippers to design and implement pragmatic, high-impact strategies that align business goals with advanced technology solutions. Tara's unique perspective blends vendor-side product leadership, hands-on implementation expertise, and operational insight—allowing her to provide objective advisory services rooted in real-world experience. Her background includes senior roles at e2open, BluJay Solutions, and LeanLogistics, where she helped shape TMS, visibility, and parcel execution capabilities for global shippers.

Rachelle Butler is a Director, Strategy at JBF Consulting with more than 15 years of experience spanning logistics operations and technology. She partners with shippers to assess, design, and implement solutions that align operational needs with long-term business direction. Rachelle's background includes roles in product leadership, consulting, implementation, and post-deployment client success at e2open, BluJay Solutions, and LeanLogistics. She began her career in the United States Marine Corps, where she gained foundational experience in transportation coordination and logistics operations. Rachelle brings a practical, real-work approach to helping clients realize meaningful value from their operational investments.

About JBF Consulting

JBF Consulting is a leading logistics strategy advisory and technology integration firm that partners with shippers to transform their logistics and supply chain execution operations. We empower clients to achieve operational efficiency and scalable, sustainable value through strategy development, roadmap orchestration, unbiased technology selection, expert implementation, data-driven insights, and ongoing managed services. For over two decades, our client-centric approach and alliances with best-of-breed solution providers have ensured that every strategy and solution we deliver drives measurable impact, long-term success, and customer satisfaction.

Contact JBF Consulting to discover how our team can help you modernize, mitigate risk, and drive sustainable growth through your transportation management system investment.

Email us at contact@jbf-consulting.com
or visit us at jbf-consulting.com.

JBF
design to deliver