

Timly

The State of AI in Inventory Management 2027

Why the smartest companies use AI and keep humans in the loop

Photo Courtesy

Table of Contents

Executive Summary	1
The Current State: Where AI Stands in Inventory and Asset Management Today	2
The Automation Trap: Where Full Automation Becomes Risky	5
The Case for Human-in-the-Loop	9
Looking Ahead: The Next 12–24 Months	13
Recommendations for Action	16
Practical Example: How Timly AI Combines Automation With Human Checkpoints	18
About Timly and References	20



Photo Courtesy

Executive Summary

AI adoption in business is growing rapidly, but a high degree of automation does not automatically indicate success. What matters is how precisely automation is applied and where it is used. This is the central thesis of this white paper, and it can be summed up in one sentence:

> *The question that determines success in asset management is not "human or machine?" but exactly where the human checkpoint sits.*

Three observations lead to this conclusion:

1. Current analyst data reveals a **significant gap** between broad AI adoption and actual scaling: In the supply chain, AI use remains largely selective. According to Gartner, only 17% of organizations are pursuing an immediate, transformative redesign of their processes using AI, while 83% are taking a more incremental approach.

2. Full automation becomes particularly risky when automated decisions have an **immediate financial or balance-sheet impact** without a human checkpoint. A case involving a related asset class illustrates these risks: Zillow discontinued its algorithm-driven home-buying program in 2021 after recording a $407.9 million inventory write-down, highlighting the financial consequences that can arise when model-driven purchasing decisions are not adequately controlled.

3. The solution is not a return to more manual work, but a deliberate architectural choice: **controlled automation** with clearly defined approval checkpoints for high-risk, difficult-to-reverse decisions, alongside a high degree of automation for low-risk, reversible decisions. A documented case from the retail sector shows how this can work in practice: At Zara, an optimization model calculated shipments to more than 1,500 stores, while a team of around 60 **employees retained the ability to override any recommendation**. According to the case study, they rarely did, but having that option was what made the system acceptable to the team.

For decision-makers in organizations with significant inventory and asset management needs, the key question is not whether to automate by using AI, but rather which decisions should even be automated, and which require a human in the loop.

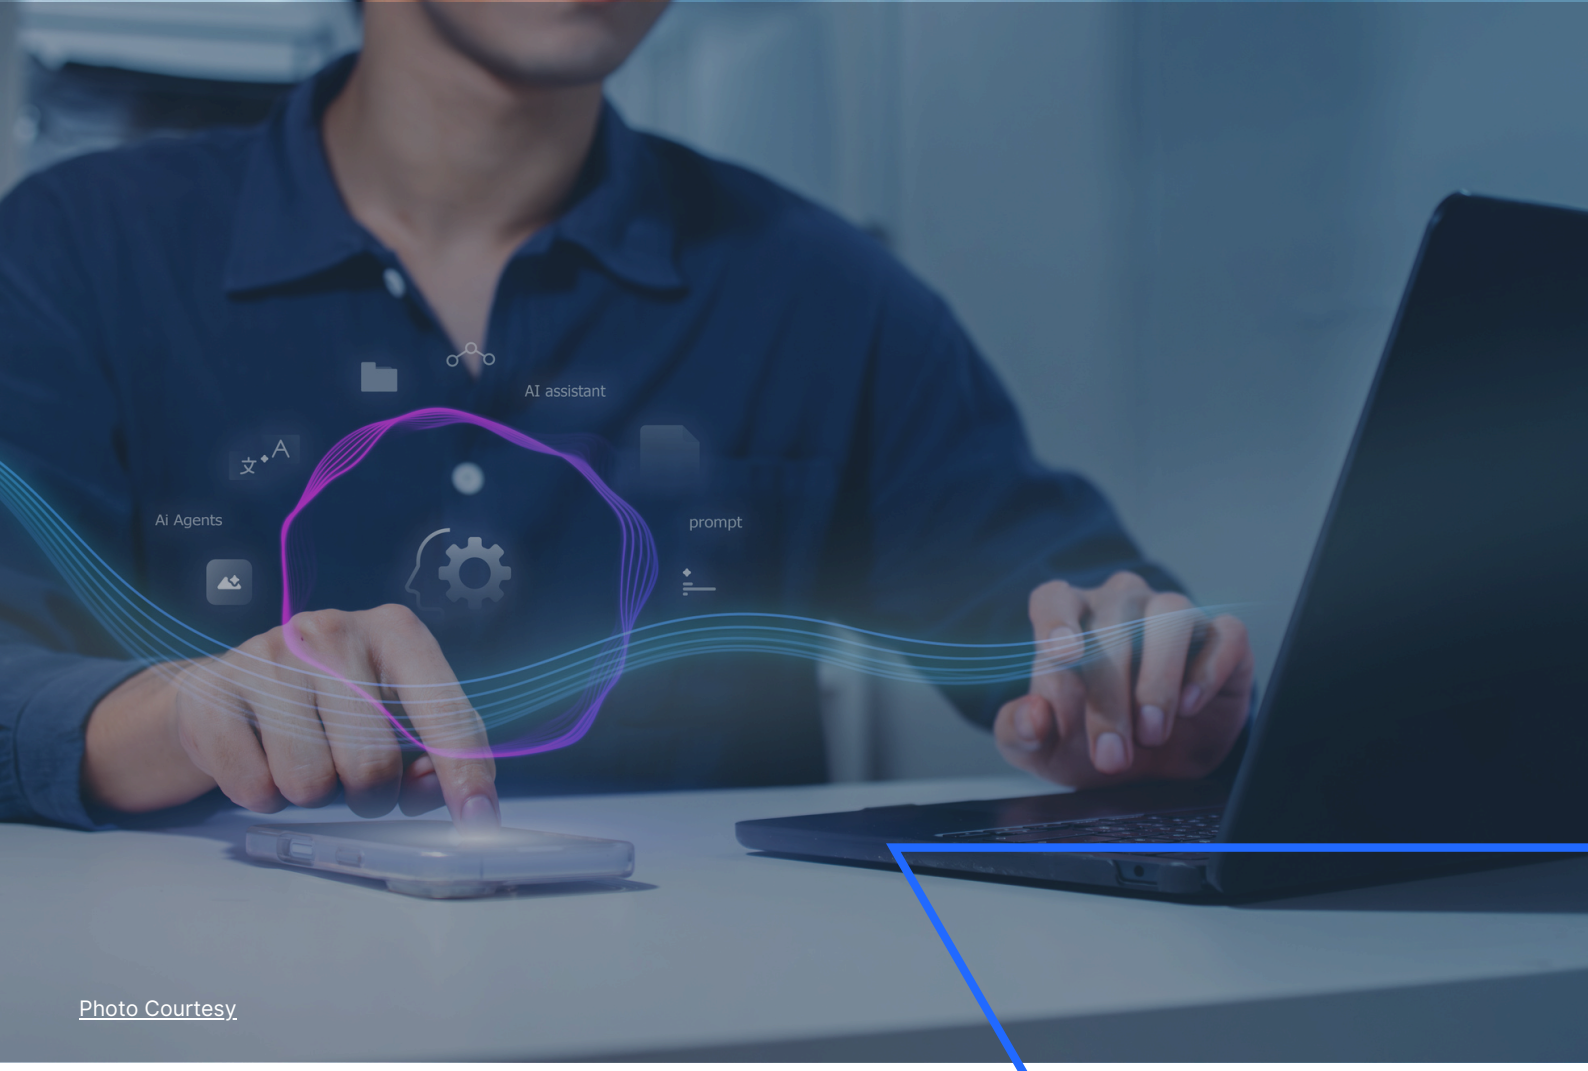


Photo Courtesy

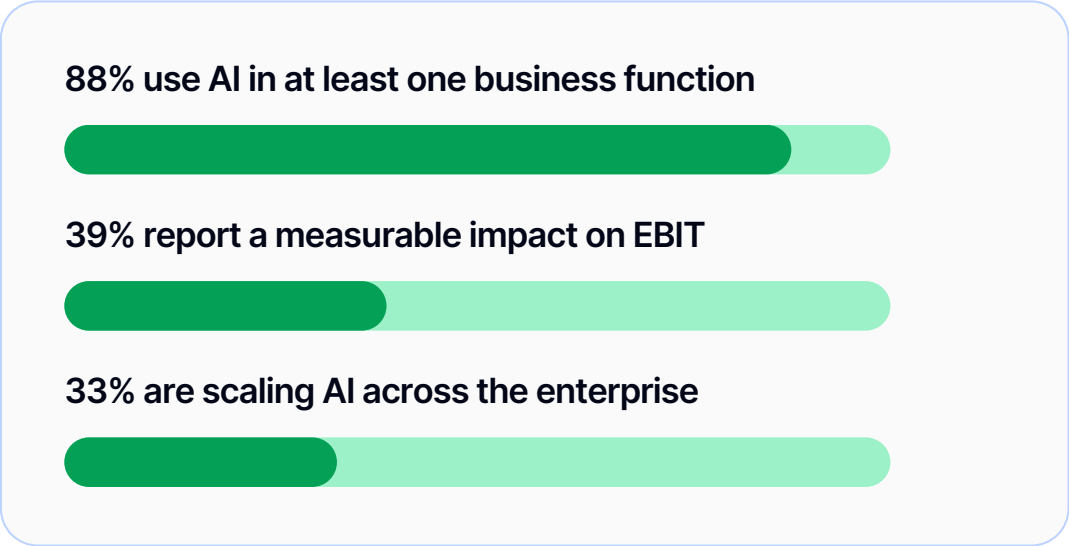
The Current State: Where AI Stands in Inventory and Asset Management Today

The Current State: Where AI Stands in Inventory and Asset Management Today

High AI adoption rates and a high level of operational maturity are two different things. Many organizations already use AI in specific areas. That alone, however, does not tell us whether these applications have moved beyond the pilot stage, scaled across the enterprise, or delivered a measurable business impact. Current analyst data helps illustrate this **gap between adoption and maturity**.

Broad Adoption, Limited Maturity

McKinsey's 2025 State of AI report highlights this distinction across industries: 88% of organizations now use AI in at least one business function, up from 78% the previous year, but **most are still in the experimentation or pilot phase**. Only about one-third of organizations analyzed have begun scaling AI programs across the enterprise, and just 39% report a measurable impact on EBIT. [1] These figures illustrate the gap between widespread AI adoption, actual scaling, and value realization.



A separate Gartner survey of 140 supply chain leaders in November 2025 shows the same distinction even more clearly in the supply chain sector: Only 17% of organizations are pursuing an immediate, transformative redesign of their processes using AI. The other 83% are either applying AI **to individual use cases** or gradually scaling it within existing processes.[2]

Investment volume itself is also an unreliable indicator of maturity. According to a Gartner survey of 394 supply chain professionals conducted between November 2025 and February 2026, supply chain organizations now allocate 67% of their digital investments to AI initiatives. At the same time, 55% of the chief supply chain officers surveyed say they **cannot clearly quantify the return on their AI investments**.[3] A strong willingness to invest therefore does not necessarily mean that the actual benefits are clear.

Three Use Cases

Within inventory and asset management, individual use cases differ significantly, both in how widely they are adopted and in the quality of the available evidence.

The Current State: Where AI Stands in Inventory and Asset Management Today

Demand Forecasting / Predictive Inventory

Demand and inventory forecasting is one of the most established AI applications in supply chain management. AI models analyze historical sales and inventory data, along with other influencing factors, to support demand forecasts as well as the inventory decisions that are based on these forecasts.

Automated Categorization and Data Capture

Image recognition for automated asset capture is gaining traction but is still much less widespread. In a December 2023 survey of 506 supply chain professionals with operational responsibilities, Gartner found that 20% of respondents had already implemented AI-enabled vision systems at that time. Gartner separately forecast that around 50% of all companies operating warehouses would use such systems by 2027. This forecast covers a different and broader population than the 20% figure, so the two values should not be treated as directly comparable.[4]

More recent data confirms the trend: Roboflow's Vision AI Trends 2026 report, published in February 2026, analyzed more than 200,000 computer vision projects implemented on its platform. Inventory and item tracking was the most common workflow among the projects analyzed, accounting for 30%.[5]

Because this analysis is based on the customer base of a single provider rather than an independent survey sample, it cannot be used to calculate an exact market share. However, the trend is consistent with Gartner's forecast.

Anomaly Detection

AI can also be used to identify shrinkage, incorrect transactions, and unusual usage patterns. These systems can then support manual inventory counts and control processes, but they do not fundamentally replace them.



Key Takeaway:

The available data shows that AI adoption and operational maturity can differ considerably, and that maturity varies by use case.



Photo Courtesy

The Automation Trap: Where Full Automation Becomes Risky

The Automation Trap: Where Full Automation Becomes Risky

Not all automated decisions carry the same level of risk. In asset and inventory management, it can be tempting to view full automation as the logical next step after successful pilot projects: If an algorithm reliably generates purchase recommendations, why not let it place the orders, too? If a model forecasts changes in value, why not have it automatically update the valuation in the financial statements?

> *Additional risks arise precisely at this transition from decision support to autonomous decision-making:* ***An automated output can have immediate operational consequences*** *without anyone checking its impact first.*

Risk Area 1: Incorrect Inventory Decisions Caused by Unsupervised Automation

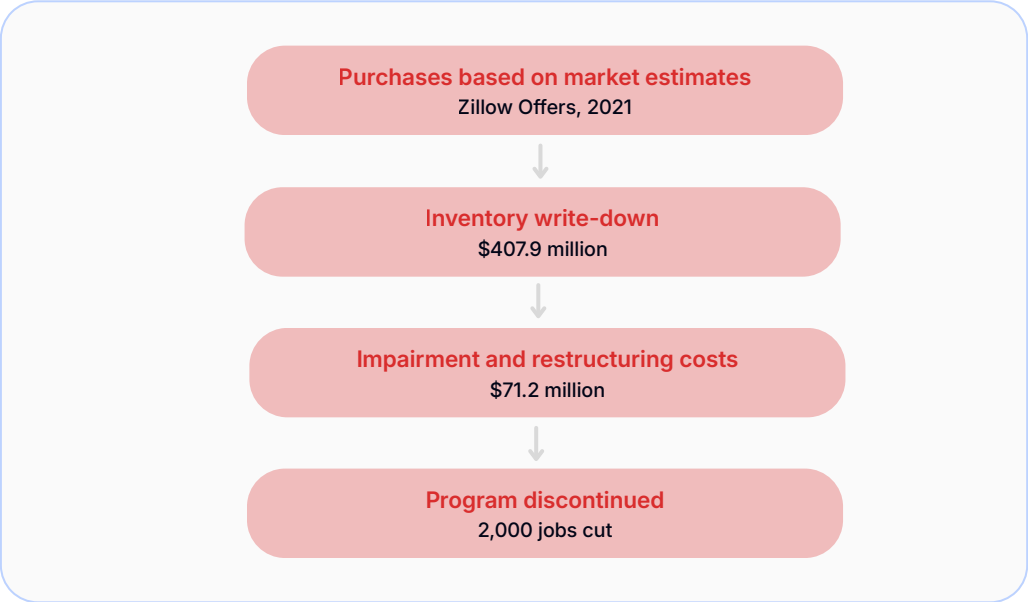
A Deloitte survey from September 2025 highlights a structural obstacle: 35% of the AI leaders surveyed identified infrastructure integration as their biggest challenge, particularly because many companies rely on rigid legacy infrastructure that is difficult to connect to automated systems, especially autonomous AI systems.[6] A separate Gartner survey from April 2026 confirms this from a supply chain perspective: 56% of chief supply chain officers identify legacy integration as one of their biggest challenges, while 50% report limited in-house expertise.[7] If companies introduce fully automated replenishment without plausibility checks on this uncertain technical and organizational foundation, data inconsistencies and system disconnects can directly affect real inventory decisions.

Risk Area 2: Incorrect Valuations Without Human Review: Financial Risk and Audit Implications

A well-documented case from a related asset class (real estate) shows how costly an automated valuation decision can become without an adequate human checkpoint. Zillow, the US real estate marketplace operator, used its algorithmic valuation model, “Zestimate,” as the basis for purchase offers through its Zillow Offers program.[8]

The company’s filings directly substantiate the following: In the third quarter of 2021, Zillow purchased numerous homes at prices above the then-current estimates of the prices at which they could later be sold. According to its 2021 annual report (Form 10-K), the company had to record a $407.9 million inventory write-down, along with $71.2 million in impairment and restructuring costs associated with winding down the business.[9] The program was discontinued entirely, and the company eliminated around a quarter of its workforce—about 2,000 jobs.[10]

The Automation Trap: Where Full Automation Becomes Risky



A more cautious but well-supported interpretation follows: The case illustrates the financial risks associated with model-driven purchasing decisions. When algorithmic estimates inform significant operational commitments without sufficient human review, **inaccurate assumptions can translate into substantial financial losses and balance-sheet implications**.

The case is also relevant regarding financial auditing. In the 2021 annual report, the auditors describe how the net realizable value of the real estate inventory was determined, including sample testing of individual homes and comparisons with actual selling prices after the balance sheet date. Because the valuation required judgment, the auditors classified it as a Critical Audit Matter and determined that **additional audit effort was required**.

In asset management, automated valuation recommendations or value-related accounting entries without plausibility checks carry a structurally similar risk, with corresponding implications for financial audits and audit traceability.

Risk Area 3: Diffusion of Responsibility

When an automated system makes a consequential decision, questions of ownership and accountability arise: Who is responsible: the system, the model's developers, or the person who failed to review the result? A June 2026 study by the IBM Institute for Business Value, based on a survey of 2,000 CIOs and CTOs worldwide, shows that this question is not always clearly resolved in practice. Two-thirds of the technology leaders surveyed say they are held accountable for AI systems they do not fully control; 70% also report that business departments adopt new technologies faster than IT can keep track of them.[11]

This demonstrates a measurable **gap between accountability and control**, not merely a theoretical risk. It also shows that responsibilities and control mechanisms must keep pace with the growing speed of AI adoption.

The Automation Trap: Where Full Automation Becomes Risky

Key Takeaway:

Many of these risks do not arise from implementing automation itself, but from situations in which automated decisions take effect immediately without a clearly defined (human) checkpoint. The Zillow case, legacy-system integration challenges, and the measured accountability gap in autonomous AI systems illustrate this from three different angles.

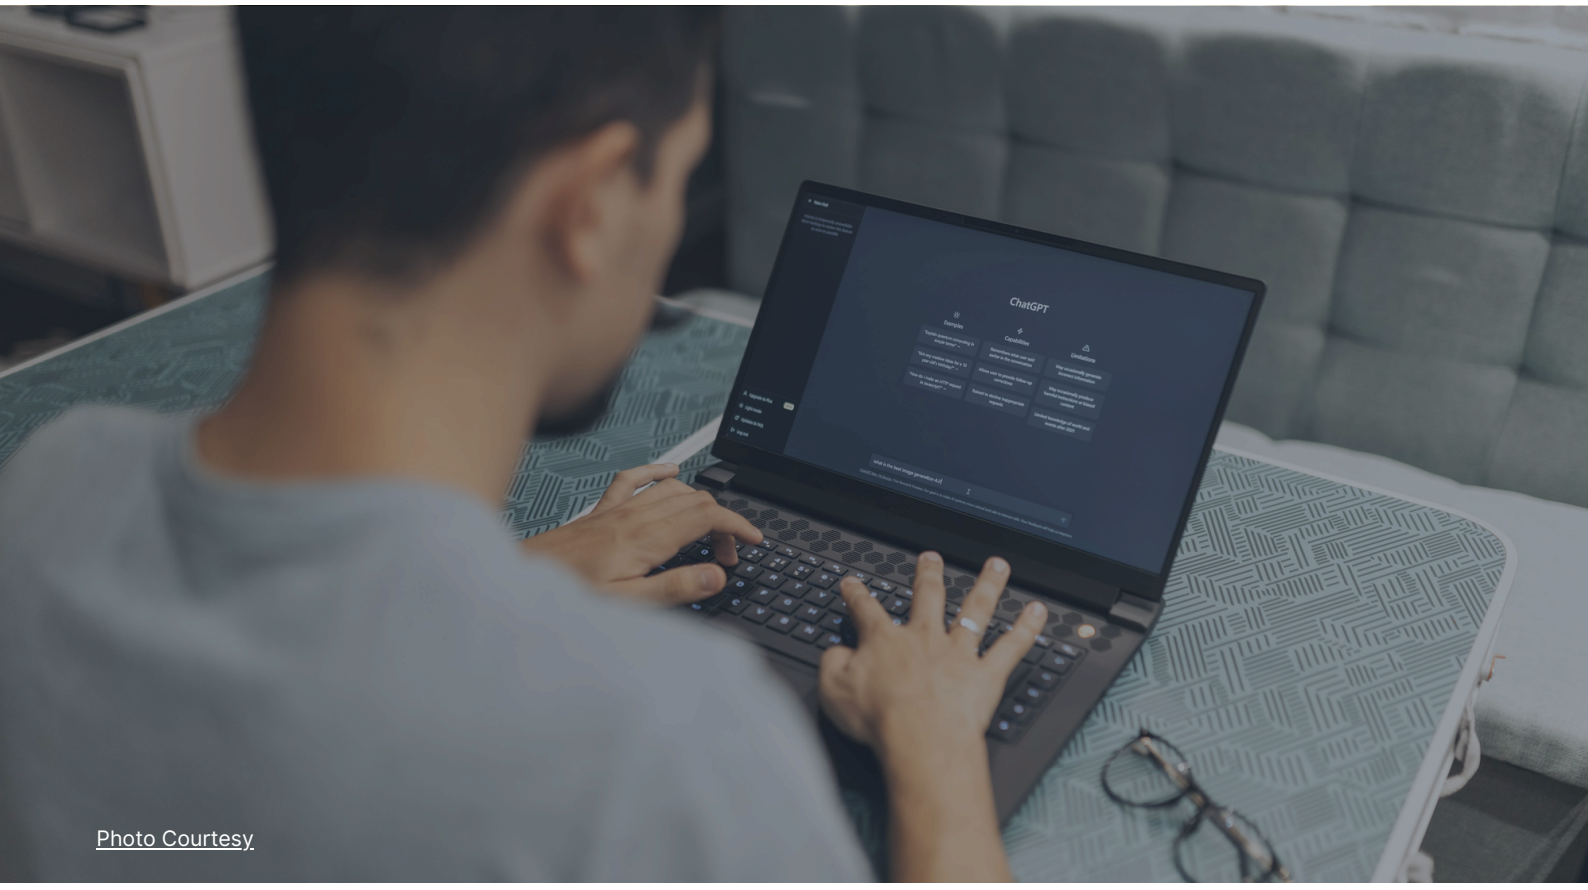


Photo Courtesy



Photo Courtesy

The Case for Human-in-the-Loop

The Case for Human-in-the-Loop

As companies delegate more and more operational decisions to systems, the question of **who can intervene when necessary** becomes increasingly important. Today, 60% of executives regularly use AI to support decisions, and according to a Gartner forecast cited by Deloitte, AI agents are expected to support or automate half of all business decisions by 2027.[12] As The Automation Trap: Where Full Automation Becomes Risky showed, increasing automation can create a gap between the speed of decision-making and the ability to maintain control. A Deloitte survey adds another perspective: 46% of companies identify governance as one of their biggest AI risks, but only 21% already have a mature governance model in place.[13]

Against this backdrop, **human-in-the-loop is not a step backward from full automation**. It is a deliberate architectural decision: When a system can make a decision quickly and reliably, it should. But when a decision has significant consequences or is difficult to reverse, at least one person must remain involved in the process.

A Documented Case: Zara

A case study by MIT Sloan and UCLA Anderson shows how this principle works in a real-world system at Zara. **Starting in 2006, an optimization model calculated weekly shipments to more than 1,500 stores worldwide.[14]** Most individual inventory decisions are low-risk and can be corrected within a few days, making them a classic use case for automation. However, Zara deliberately retained the ability for its approximately 60-person warehouse team to manually override any recommendation before shipment.

Notably, this override function was rarely used once the system was in operation. After a short testing period, employees almost never changed the model's recommendations manually. What mattered, however, was that they always had the option to do so. This freedom to override the system's output gave the team the confidence to embrace the new system in the first place. The checkpoint did not need to be used constantly to serve its purpose. Its mere existence was enough.

What makes the case instructive is that Zara kept an override even for low-risk, reversible decisions. **Not every decision needed review, but having the option made the system acceptable** to the people working with it.

A Simple Decision-Making Principle

For practical purposes, a framework with two dimensions is enough: the risk of making the wrong decision and its reversibility (how easily a mistake can be undone).

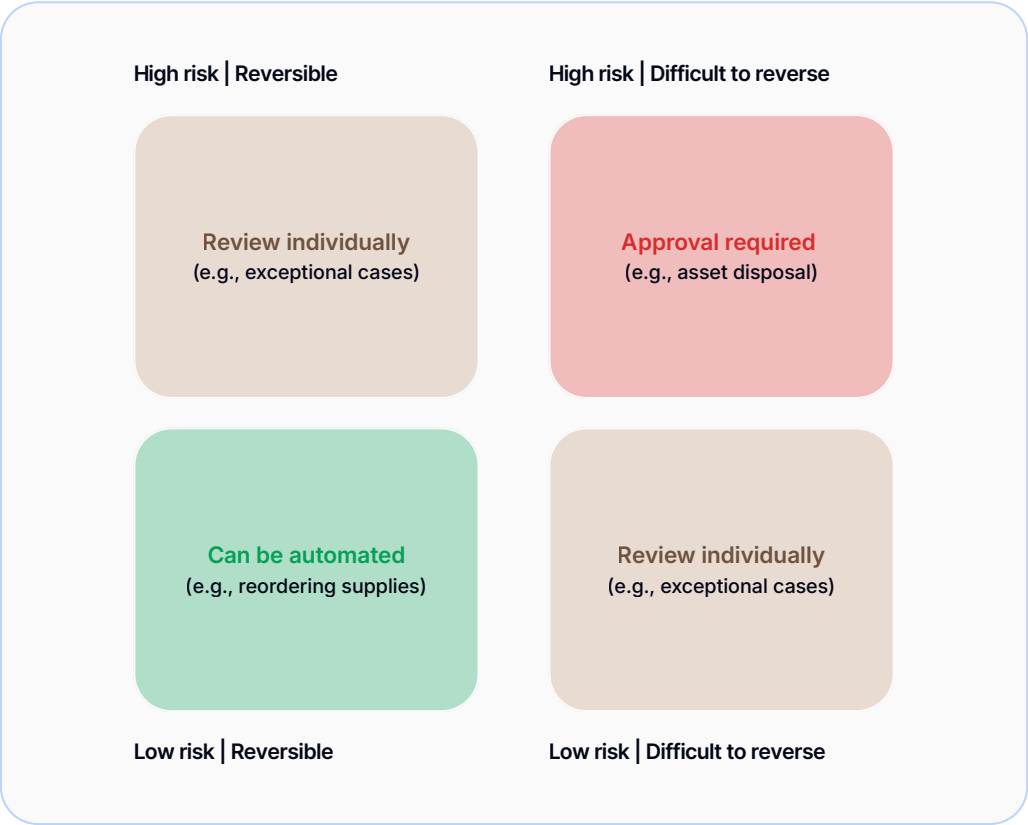
- **Low risk, easily reversible**

  Example: Reordering consumable supplies. If too much or too little is ordered, the issue can be corrected in the next ordering cycle. Decisions like these can be fully automated.

- **High risk, difficult to reverse**

  Example: Disposing of a fixed asset or recording an impairment. A fixed asset that has been incorrectly written off cannot simply be brought back, and an incorrect impairment affects the balance sheet and audit results. Decisions like these require approval.

The Case for Human-in-the-Loop



This principle can be applied directly to common asset management scenarios in which approval processes make sense:

- Fixed asset additions and disposals: physical and accounting consequences, often difficult to reverse
- Impairments and depreciation: direct effects on financial statements and tax liabilities
- Decisions that must be included in an audit trail: anything that must be fully traceable during a financial audit

Why a Checkpoint Must Be More Than a Formality

For consequential applications, review by human experts still remains the most effective way to prevent critical errors.[15] However, this review must be genuine, not merely procedural. People who supervise systems that have performed reliably over a long period tend to accept their results without checking them and gradually relinquish their own judgment. This phenomenon is known as **automation bias**.[16] The **EU AI Act** explicitly addresses this issue: People responsible for oversight must be made aware of the tendency to over-rely on AI results.[17] This regulatory framework is examined in more detail in Looking Ahead: The Next 12–24 Months.

The Case for Human-in-the-Loop

The Zara case shows one part of the answer: control that can be exercised at any time, and whose availability is known to everyone involved. But an override that is rarely used is also where automation bias thrives. For consequential decisions, the **checkpoint needs active safeguards**. These include showing reviewers the data behind each recommendation, requiring a brief justification for approvals, and regularly sampling approved decisions for review.

For IT infrastructure workflows, the share of decisions fully approved by humans is expected to fall from 95% to 40% by 2028, meaning autonomy is increasing while approval processes are shrinking. [18] This may be acceptable for highly standardized IT operations. Asset management is different, though: Individual mistakes can have accounting and legal consequences. The goal should not be to increase the number of approvals, but to make them more targeted. Not every decision needs a person, but **every consequential decision needs a clearly defined, genuine checkpoint**.

> The question is not “human or machine?” but “Where exactly does the human checkpoint sit?”



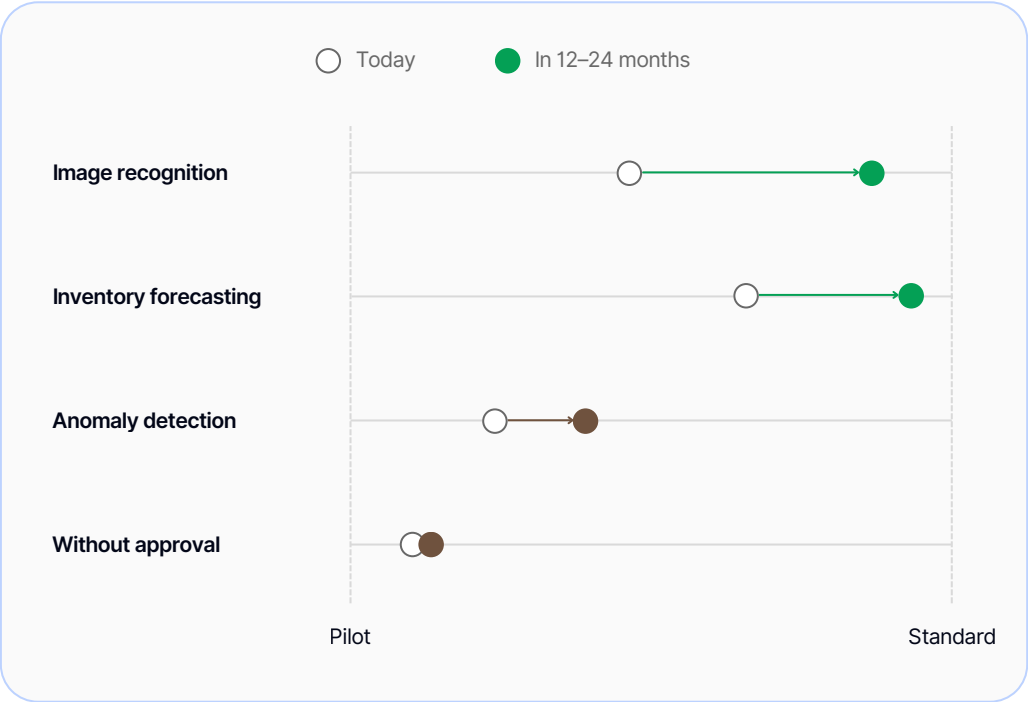
Photo Courtesy



Photo Courtesy

Looking Ahead: The Next 12–24 Months

Looking Ahead: The Next 12–24 Months



Automated Categorization and Image Recognition Will Become Standard

The technology is maturing quickly. Supply chain leaders in healthcare are already being encouraged to gradually replace manual inventory counting with ceiling-mounted computer vision systems and predictive AI.[19] Such systems build on barcodes and RFID rather than replacing them. **Reliable identification remains the foundation**, while image recognition takes over more of the manual capture work. Together with the Roboflow data cited earlier, this suggests that image recognition in asset capture will move from pilot projects into standard operations over the next 12–24 months.

Demand Forecasting / Predictive Inventory Will Become More Widespread

Among leading supply chain organizations, adoption of AI-powered inventory forecasting is already at 87%, with forecast accuracy improving by 20–40%.[20] Walmart has publicly stated, among other things, that its "Self-Healing Inventory" system has saved more than $55 million and is now operating across several countries.[21] Over the next 12–24 months, this use case is likely to evolve from a competitive advantage for market leaders into a standard feature of ERP environments at midsized companies as well.

Anomaly Detection Will Grow, but Remain a Supporting Tool

For detecting shrinkage, incorrect transactions, and unusual usage patterns, adjacent fields such as fraud detection offer a model: **AI-based pre-screening combined with human review is the established approach** for keeping false positives manageable. The value here lies in filtering and prioritization, not autonomous decision-making. For the foreseeable future, this use case will remain assistive rather than autonomous.

Looking Ahead: The Next 12–24 Months

Fully Automated Inventory Decisions Without Approval Will Remain a Niche

By the end of 2027, more than 40% of agentic AI projects are expected to be canceled because of unclear business value, rising costs, or inadequate risk controls.[22] Only 5% of the generative AI pilots examined show a measurable impact on business results.[23] Boston Consulting Group also notes that the transition to an AI-first operating model succeeds only when governance is embedded in the operating architecture from the outset, ensuring auditability, model traceability, and clear accountability.[24] Taken together, these findings suggest that **fully automated decisions without such a governance foundation will remain the exception rather than the rule**.

Regulation and Standards Will Guide the Way, Not Just Impose Requirements

Traditional inventory management algorithms are not listed as high-risk use cases in Annex III of the EU AI Act and therefore generally fall into the lowest risk category; as a result, they do not trigger the Act's most stringent obligations. [25] Nevertheless, the regulation shapes expectations more broadly: High-risk systems are subject to a requirement for human oversight. The transition period originally scheduled for August 2, 2026, was postponed to **December 2, 2027**, by the "Digital Omnibus on AI" (Regulation (EU) 2026/1744), which entered into force in July 2026. [26] This is a reminder that the regulatory framework itself is still evolving. At the same time, ISO/IEC 42001 has established the first certifiable, globally applicable standard for AI management systems, explicitly identifying human oversight as a control objective. [27]

For companies that are not directly subject to the EU AI Act, **ISO/IEC 42001 offers a practical, voluntary framework**. Decisions affecting financial statements and audits (e.g., impairments, depreciation, and fixed-asset disposals) are already subject to existing traceability requirements under auditing and accounting standards. Regulation and standardization therefore reinforce expectations that have long existed in asset management through audit standards, rather than creating them from scratch.

Progress will come not from more automation, but from better control architecture.



Photo Courtesy



Photo Courtesy

Recommendations for Action

Recommendations for Action

1. **Classify your asset-related decisions by risk and reversibility** before discussing the level of automation. A simple two-axis framework is enough to get started. The classification itself provides more clarity than any subsequent tool decision.

2. **Define clear approval checkpoints** for high-risk decisions that are difficult to reverse. Determine who approves them, what data they need, and within what time frame. Do so before the first system goes live, not afterward.

3. **Start with use cases that are already in production**, not the most ambitious or complicated ones. Inventory forecasting and automated categorization already deliver demonstrable value. Fully automated approval processes are not there yet.

4. **Build in audit trails right from the start**, rather than adding them later. Use established control objectives, such as those in ISO/IEC 42001, as a guide. Every automated decision should be logged in a traceable way from day one; adding traceability later is considerably more expensive.

5. **Treat AI adoption as a process issue,** not just a tool decision. The software is the easy part. Checkpoints, accountability, and approval logic must be actively designed.

Structured, controlled asset digitization therefore depends above all on a carefully designed balance between automation and checkpoints. This is where Timly supports companies that want to digitize inventory and asset management without giving up control over critical decisions.

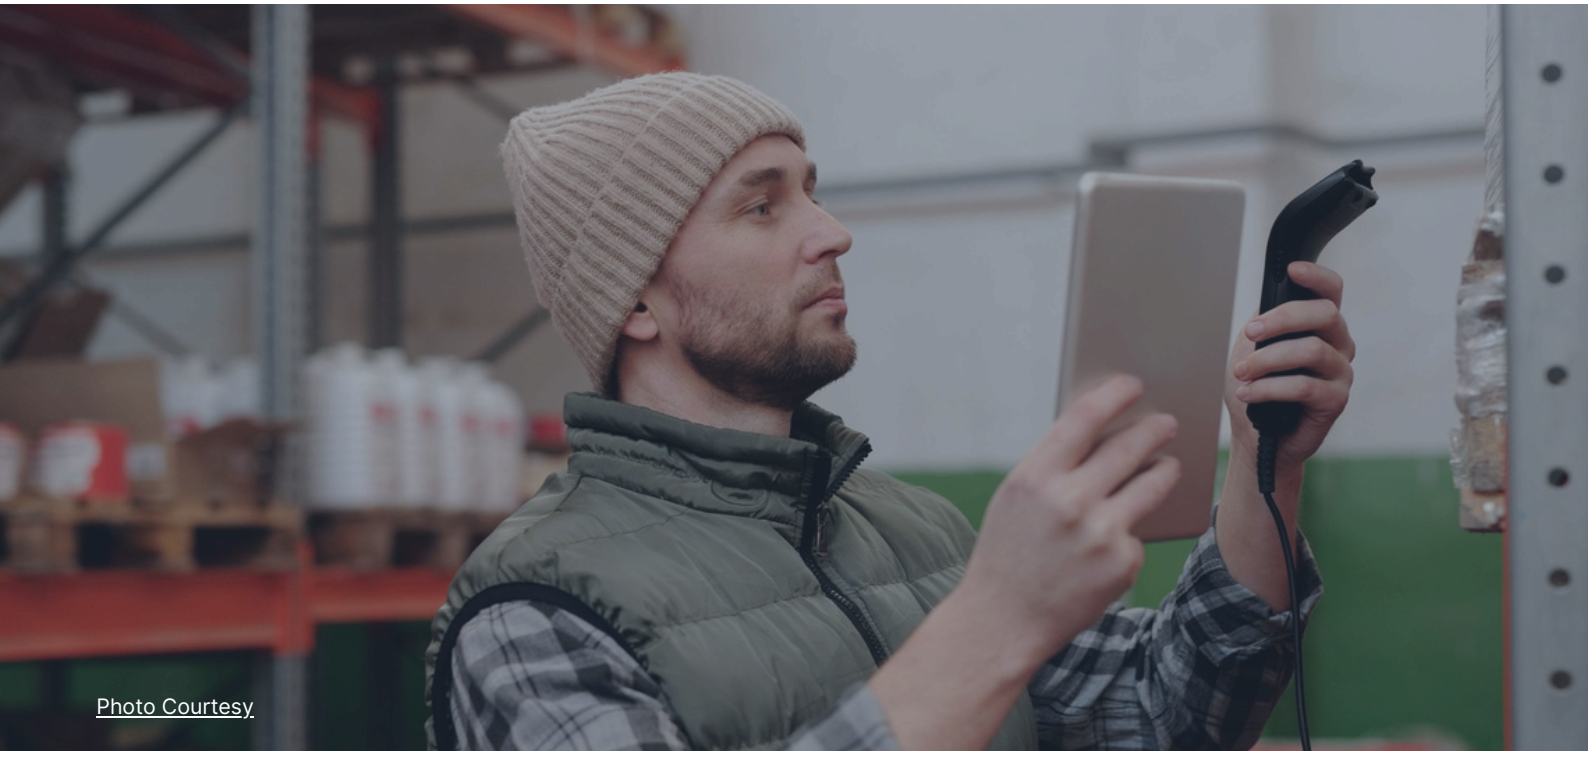
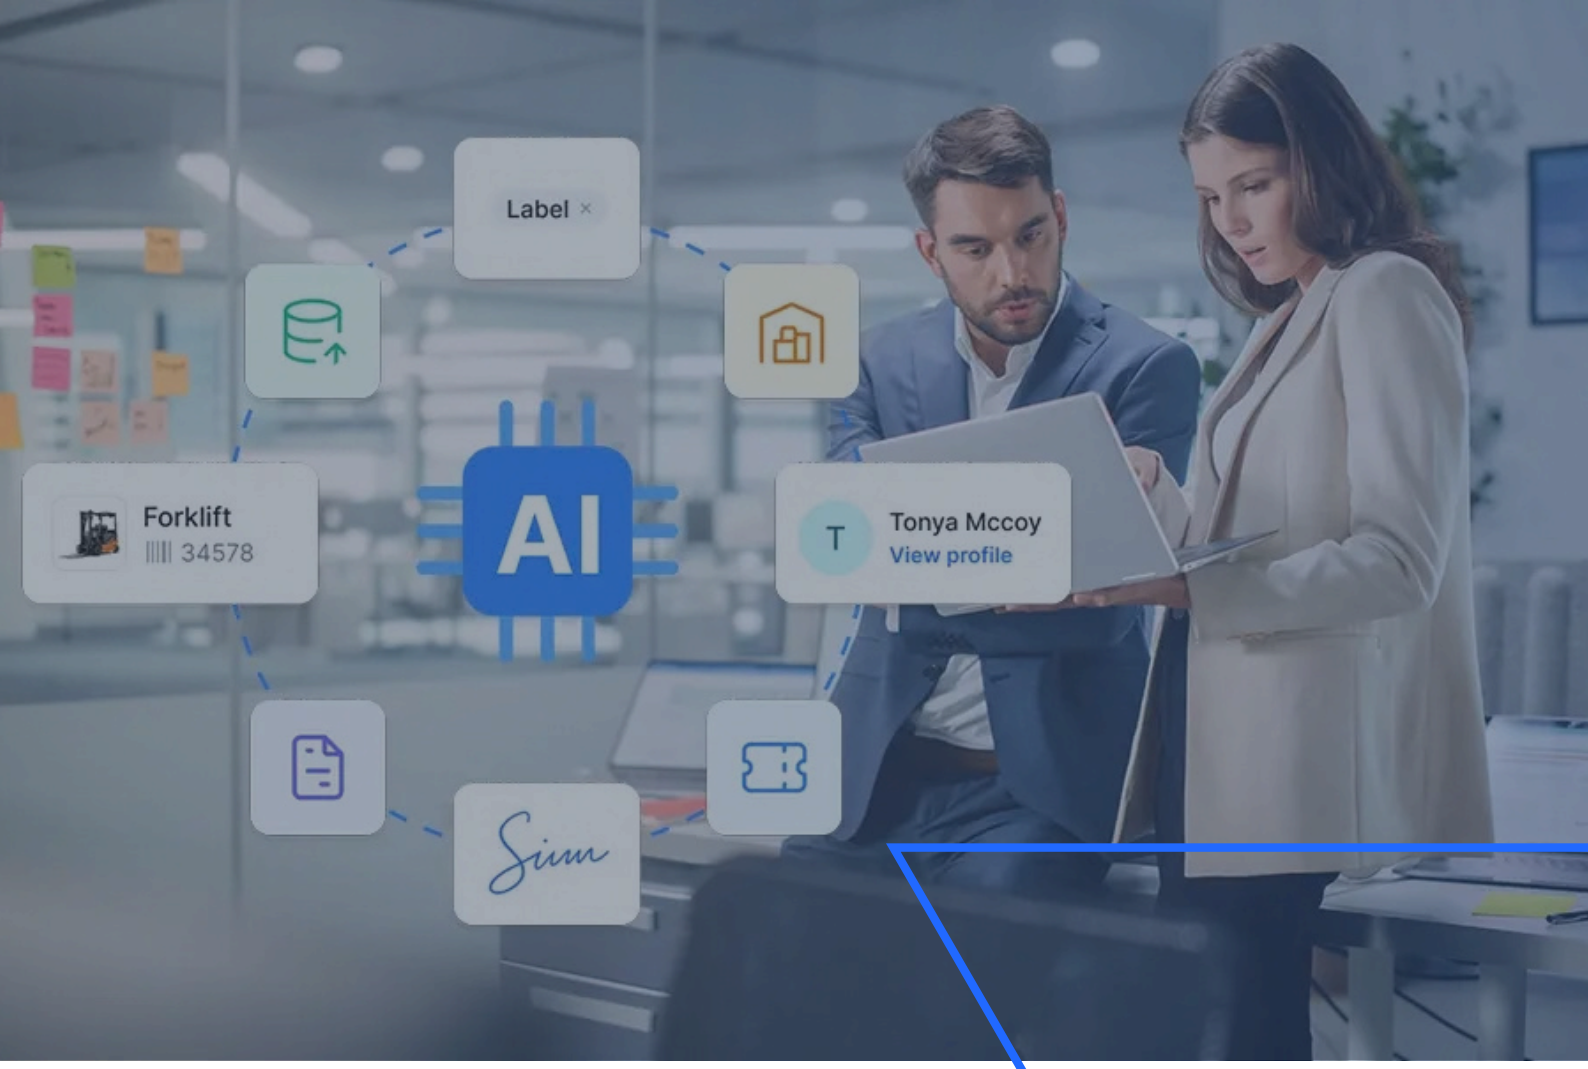


Photo Courtesy

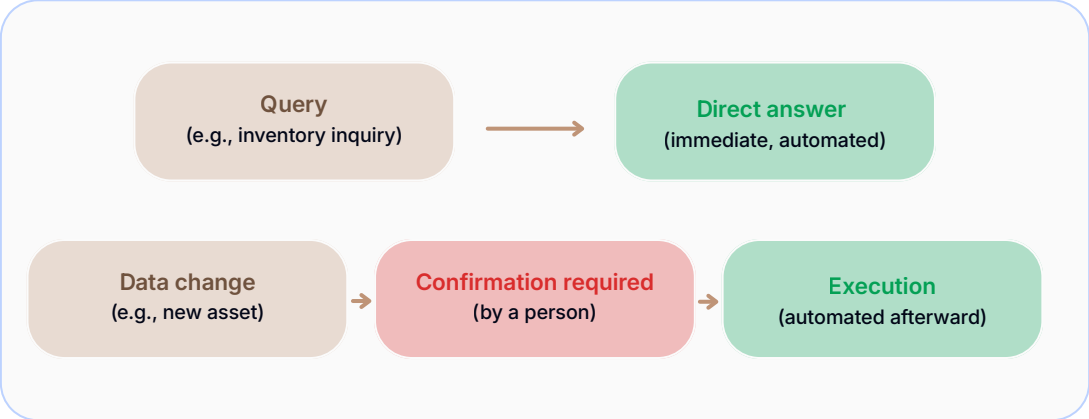


Practical Example: How Timly AI Combines Automation With Human Checkpoints

Practical Example: How Timly AI Combines Automation With Human Checkpoints

The decision-making principle described in The Case for Human-in-the-Loop can also be seen in Timly AI itself. With Timly MCP and the Timly AI Assistant, users can manage inventory through natural language: from creating locations and assets to issuing and returning them, as well as importing CSV files from external systems such as SAP or Excel.

Here, too, the key is not automation itself, but where the checkpoint sits. The AI answers queries directly, such as "What has Michael checked out right now?" or "What's in vehicle 413?" Any action that changes data, however, whether it's creating an asset, changing a status, or assigning an item, is presented to the user for confirmation before it is carried out. This is a deliberately conservative application of the two-axis framework. Reading data carries no risk and runs automatically. Every change to the inventory, even an easily reversible one, passes through a human checkpoint, so **no change takes effect without a responsible person seeing it first**.



Two additional design principles reinforce this white paper's core arguments:

- **Inherited permissions instead of separate AI permissions:** The assistant inherits exactly the permissions of the signed-in user. No more and no less. This means accountability is structurally anchored to the responsible person, not the AI itself.
- **Comprehensive audit log:** Every action, session, and tool call is logged and can be traced. This fulfills the audit-trail requirement described in Recommendations for Action from the outset.

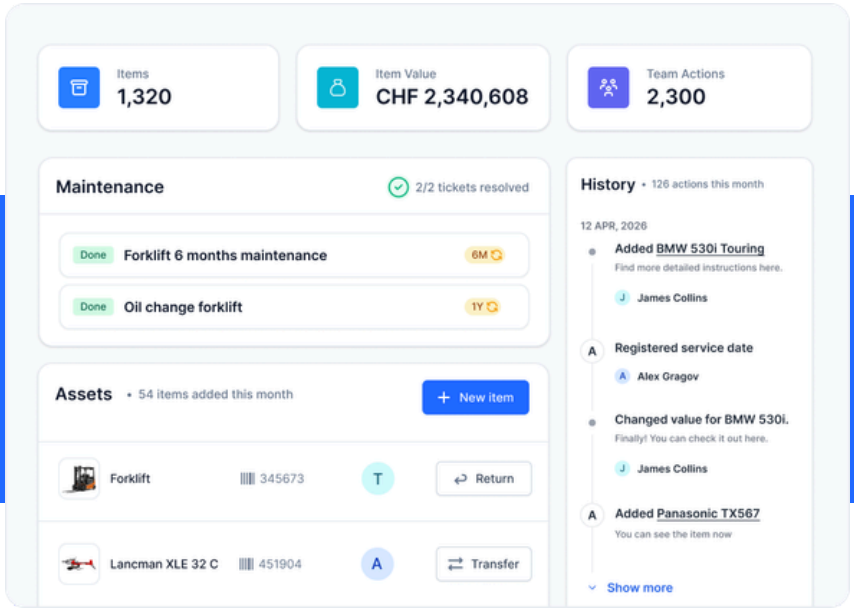
A practical example shows how this works day to day: An SAP export containing several hundred assets can be imported using copy and paste. The AI automatically identifies the columns and prepares records for people not yet in the system. The actual import into the inventory system, however, takes place only after the person responsible confirms it. This **combination of speed during preparation and a deliberate approval checkpoint** for the actual inventory change is the "human-in-the-loop principle" described in this white paper.



Photo Courtesy

About Timly and References

Timly



About Timly and References

Timly helps companies manage physical assets and equipment centrally, on the go, and with full traceability. This provides the data foundation for the kind of controlled automation described in this white paper. The platform is used particularly in industries with distributed and dynamic asset structures, including construction, manufacturing, facility management, and technical services.

A key focus is **connecting physical asset visibility with the control and traceability requirements** that become increasingly important as automation grows. Key capabilities include:

- Centralized management of physical assets throughout their life cycle
- Mobile data capture using a smartphone or tablet
- QR code- and barcode-based inventory tracking as a reliable data foundation for automated inventory forecasts
- Role-based responsibilities and defined approval checkpoints for critical decisions
- Full traceability for financial audits and audit requirements
- Support for distributed locations and mobile teams

Timly demonstrates how these principles can also be implemented technically in its own product. See Practical Example: How Timly AI Combines Automation With Human Checkpoints for a specific example using Timly MCP and the Timly AI Assistant. By combining easy data capture, mobile use, and a centralized, auditable data foundation, Timly helps companies strike the balance described in this white paper: a high degree of automation where it makes sense, and a clearly defined, traceable checkpoint for decisions with significant consequences.

Book a free demo to see how controlled automation can work in your asset management processes.

SCHEDULE A DEMO

References

[1] McKinsey, « The State of AI: Global Survey », 2025 - mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

[2] Gartner, press release, "Gartner Survey Shows AI Is Not Driving Supply Chain Operating Model Transformation," May 2026 - gartner.com/en/newsroom/press-releases/2026-05-06-gartner-survey-shows-ai-is-not-driving-supply-chain-operating-model-transformation

[3] Gartner, press release, "Gartner Survey Finds Majority of Chief Supply Chain Officers Unclear on AI Investment Returns," August 2026 - gartner.com/en/newsroom/press-releases/2026-08-05-gartner-survey-finds-majority-of-chief-supply-chain-officers-unclear-on-ai-investment-returns

[4] Gartner, press release, "Gartner Predicts Half of Companies With Warehouse Operations Will Leverage AI-Enabled Vision Systems by 2027," June 2024 - gartner.com/en/newsroom/press-releases/2024-06-12-gartner-predicts-half-of-companies-with-warehouse-operations-will-leverage-ai-enabled-vision-systems-by-2027

[5] Roboflow, "Vision AI Trends 2026," February 2026 - blog.roboflow.com/vision-ai-trends-2026/

[6] Deloitte, "AI Adoption Challenges: AI Trends" (Pulse Check Series), September 2025 - deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/blogs/pulse-check-series-latest-ai-developments/ai-adoption-challenges-ai-trends.html

[7] Gartner, press release, "Gartner Survey Finds Technology Integration and Talent Perceived as Key Roadblocks to Scaling AI in Supply Chain," April 2026 - gartner.com/en/newsroom/press-releases/2026-04-29-gartner-survey-finds-technology-integration-and-talent-perceived-as-key-roadblocks-to-scaling-ai-in-supply-chain

[8] CNN Business, "What Zillow's Exit From the iBuying Business Means for Home Values," November 2021 - edition.cnn.com/2021/11/09/tech/zillow-ibuying-home-zestimate

[9] Zillow Group, Inc., Form 10-K for fiscal year 2021, SEC filing - sec.gov/Archives/edgar/data/1617640/000161764022000013/z-20211231.htm

[10] Axios, "Zillow to Wind Down Home-Flipping Business, Cut a Quarter of Staff," November 2021 - axios.com/2021/11/02/zillow-abandon-home-flipping-algorithm

[11] IBM Institute for Business Value, "New IBM Study Finds CIOs and CTOs Face Growing AI Control Gap as Enterprise Deployment Scales," June 2026 - newsroom.ibm.com/2026-06-08-new-ibm-study-finds-cios-and-ctos-face-growing-ai-control-gap-as-enterprise-deployment-scales

[12] Deloitte, "AI and the Future of Human Decision-Making," Global Human Capital Trends 2026 (includes a cited Gartner forecast on AI agents and business decisions) - deloitte.com/us/en/insights/topics/talent/human-capital-trends/2026/decision-making-with-ai.html

[13] Deloitte, "State of AI in the Enterprise," 2026 (governance gap: 46% vs. 21%) - deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/content/state-of-ai-in-the-enterprise.html

[14] Caro, F. / Gallien, J. (MIT Sloan & UCLA Anderson), "Zara Uses Operations Research to Reengineer Its Global Distribution Process," 2009 - web.mit.edu/jgallien/www/ZaraInterfacesPaperAugust3.pdf

[15] McKinsey, "How Financial Institutions Can Improve Their Governance of Gen AI," 2025 - mckinsey.com/capabilities/risk-and-resilience/our-insights/how-financial-institutions-can-improve-their-governance-of-gen-ai

[16] Lyell, D. / Coiera, E., systematic review on automation bias, Journal of the American Medical Informatics Association, 2017

[17] EU AI Act, Article 14 (human oversight) and Article 14(4)(b) (requirement to raise awareness of automation bias) - ai-act-service-desk.ec.europa.eu/en/ai-act/article-14

[18] Gartner, "Predicts 2026: AI Agents Will Transform IT Infrastructure and Operations," Cameron Haight, Ashish Banerjee, Joe Antelmi, Jonathan Forest, Chandra Mukhyala, December 4, 2025 - itential.com/resource/analyst-report/gartner-predicts-2026-ai-agents-will-reshape-infrastructure-operations

[19] Gartner, "Getting to Inventory Autonomy: Build the Self-Managing Supply Room," August 2026 (summarized by HIT Consultant) - hitconsultant.net/2026/08/28/gartner-healthcare-csco-autonomous-supply-rooms-ai-computer-vision-inventory

[20] Supply Chain Management Review, "AI in the Supply Chain: From Pilot Programs to P&L Impact," April 2026 (includes Gartner/McKinsey figures) - scmr.com/article/ai-in-the-supply-chain-from-pilot-programs-to-pl-impact

[21] Walmart, "Walmart's U.S. Supply Chain Playbook Goes Global—and It's Reinventing Retail at Scale," July 17, 2025 - corporate.walmart.com/news/2025/07/17/walmarts-us-supply-chain-playbook-goes-global-and-its-reinventing-retail-at-scale

[22] Gartner, "Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027," June 25, 2025 - gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027

[23] MIT NANDA, "The GenAI Divide: State of AI in Business 2025," August 2025 - mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf

[24] Boston Consulting Group (Garro, Caffrey, Shetty et al.), "Supply Chain Planning 2026: Why AI Alone Isn't Enough," February 2026 - bcg.com/publications/2026/supply-chain-planning-why-ai-alone-isnt-enough

[25] Regulation (EU) 2024/1689 (EU AI Act), Article 6 and Annex III - eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689

[26] Regulation (EU) 2026/1744 ("Digital Omnibus on AI"), in force since July 27, 2026—postponing the Annex III transition deadline to December 2, 2027; summarized by Gibson Dunn, "EU AI Act Omnibus Agreement—Postponed High-Risk Deadlines and Other Key Changes" - gibsondunn.com/eu-ai-act-omnibus-agreement-postponed-high-risk-deadlines-and-other-key-changes/

[27] ISO/IEC 42001, AI management system standard with human oversight as a control objective, published in 2023 - iso.org/standard/42001