



F&B PLANT DATA MATURITY · COMPANION REPORT

# The Foundation Before the Algorithm

*What the current state of AI and analytics in food and beverage manufacturing means for your plant, and what to do next.*

---

---

Prepared by **Haystack Data Solutions** for leaders who completed the F&B Plant Data Maturity Scorecard

A food intelligence company · [haystackdata.com](https://haystackdata.com)

Start here

## How to read this report

---

*You just finished the Scorecard. This is the context behind your result.*

The Scorecard placed your plant on a five-level maturity ladder across seven categories. This report does two things with that result. First, it gives you an honest picture of where AI and analytics actually stand in food and beverage manufacturing right now, what is working, what is hype, and why so many projects stall before they pay for themselves. Second, it walks through each of the seven categories and describes what good looks like, along with the single next move for a plant at your level. The general sections apply to everyone. The tier-by-tier notes inside each category are where you find your own plant.

One idea runs through all of it. You cannot build reliable analytics or AI on a broken data foundation. Build the foundation first, and every step after it gets easier and cheaper. Skip it, and you join the majority whose pilots never reach production. Most food and beverage manufacturers use less than half the data they collect, and more than four out of five AI projects fail, usually because the data underneath them was not ready. Your Scorecard is a map of where that is true for you and where it is not.



A modern F&B production line. The equipment varies by product, bottles, cans, dough, but the data problem underneath it does not.

The short version

- **Most F&B plants are data-rich and decision-poor.** The data exists. It just cannot move, cannot be trusted, and cannot be compared across lines. Your Scorecard shows where that is true for you.
- **The value is real when the order is right.** Nestlé, PepsiCo, and General Mills are posting documented, sourced results. None of them started with AI. They started with connected, governed data.
- **Sequence beats speed.** Move up the ladder in order: foundation, connectivity, visibility, analytics, then AI. FSMA 204's traceability deadline is now July 20, 2028, which buys time to build the data plumbing that compliance and AI both depend on.




Figure 1. The five-level plant data maturity ladder. Most F&B plants cluster at Levels 2 and 3. The goal is not the label. It is the next rung.

What is inside

|                                                         |    |
|---------------------------------------------------------|----|
| Data-rich, decision-poor                                | 03 |
| The current state of AI and analytics in F&B            | 04 |
| The AI value curve: why value dips before it rises      | 05 |
| The seven categories, with your next move at each level | 06 |
| Documented results, the cost of inaction, and FSMA 204  | 11 |
| Sequencing, recommendations, and sources                | 14 |

## Data-rich, decision-poor

*Most plants do not have a shortage of data. They have data that does not add up to a decision.*

Walk any modern food or beverage plant, a bottling line, a bakery, a cannery, a snack line, and you will find data everywhere. PLCs on the mixers, checkweighers on the lines, a SCADA historian somewhere, an ERP that finance trusts, a QMS the food safety team lives in, and a dozen spreadsheets that hold the operation together. The problem is almost never a shortage of data. The problem is that the data does not add up to a decision.

This is the data-rich, decision-poor reality, and the numbers back it up. In Rockwell Automation's 2025 State of Smart Manufacturing Report, food and beverage organizations reported that less than half the data they collect, on average 44.97%, is actually used effectively. A separate manufacturing study found that only 9% of manufacturers can find the root cause of a shop floor issue immediately, and roughly half still rely on manual logs, paper, and spreadsheets even as sensors multiply. So the plant generates more data every year and gets no closer to an answer. The information exists. It just cannot move, cannot be trusted, and cannot be compared across lines or shifts.

That is the context for the wave of AI enthusiasm hitting the food industry, and it is why so much of the conversation is unhelpful. The market forecasts are enormous and the vendor decks are breathless. What gets lost is a simple fact: you cannot build reliable analytics or AI on a broken data foundation. When the underlying data is fragmented, inconsistent, or locked in silos, an AI model does not fix that. It inherits it, and then it produces confident nonsense that your operators correctly refuse to trust.

RAND Corporation's 2024 study, built on interviews with 65 data scientists and engineers, found that more than 80% of AI projects fail, roughly twice the failure rate of IT projects that do not involve AI. The reasons are not exotic, and they are not about the math. Fully 84% of the industry interviewees pointed to leadership-driven issues rather than the algorithm itself: unclear problem definition, weak sponsorship, and, above all, data that was not ready. Informatica's 2025 survey of data leaders put data quality and readiness at the top of the obstacle list.

*Foundation before AI. Build the data foundation, connect the systems, make performance visible, layer in analytics, and only then reach for AI.*

*Done in that order, each step pays for the next. Done out of order, you get a pilot that demos well and never ships.*


Your Scorecard result places you on a five-level ladder across seven categories, from Level 1, the Reactive Plant, through Level 2 Connected, Level 3 Centralized, Level 4 Insight-Driven, to Level 5, the Intelligent Plant. Most of the industry is clustered in the low to middle tiers, and that is not a reason for concern. It is just where the work is. The point of the ladder is to show the next rung and the specific move that gets you there.

## AI and analytics in F&B, told straight

*Adoption is real, uneven, and mostly early.*

Start with intent. In Rockwell's 2025 survey of more than 1,500 manufacturers, 95% said they have invested in or plan to invest in AI or machine learning over the next five years. Among food and beverage organizations specifically, 96% have invested or plan to invest in generative or causal AI and 93% in machine learning. Quality control is the top AI use case, cited by 50% of manufacturers for 2025. The appetite is not in question.

Now the reality gap. McKinsey's 2024 survey of consumer goods leaders found 71% had adopted AI in at least one business function, up from 42% a year earlier, with 56% regularly using generative AI. But McKinsey was equally clear that no consumer goods player has truly scaled its AI capabilities, and that the sector lags others in adoption depth. The pattern across sources is pilot rich, production poor. One widely cited MIT study in 2025 found that 95% of generative AI pilots delivered no measurable financial return. S&P Global found that 42% of companies abandoned most of their AI initiatives in 2025, up from 17% the year before. None of this is a reason to avoid AI. It is a reason to refuse to start with it.



**Figure 2.** The intent-to-impact gap. Nearly everyone plans to invest. Most pilots never turn into production value. Sources: Rockwell 2025, McKinsey 2024, MIT 2025, RAND 2024.

There is real movement, though. Analyzing food and beverage responses, Consumer Goods Technology reported that across manufacturing, the share of companies pushing AI beyond pilots more than tripled between 2024 and 2025, rising from 4% to 14%, and that 16% of F&B manufacturers now say more than half of their AI pilots have reached scale across sites. That is

the leading edge pulling away from the middle market. The plants that built the foundation first are converting pilots to production, and the gap between them and everyone else is widening.

The leaders are not winning because they bought better algorithms. They are winning because, years ago, they did the unglamorous work: they consolidated data, standardized it, governed it, and made it flow. When generative AI arrived, they had a clean, connected base to point it at. Everyone else is now discovering that step cannot be skipped.


For a mid-market food or beverage manufacturer, that is good news. You are not behind because you lack a data science team or a generative AI copilot. The highest-return work available to you right now is foundational, and it is work you can start without a single AI model.

## Why value dips before it rises

*There is a documented reason early digital investment can feel like it is not working.*

Economists Erik Brynjolfsson, Daniel Rock, and Chad Syverson call it the Productivity J-Curve. When a general-purpose technology such as AI enters an organization, measured productivity dips before it rises. The dip happens because the organization is pouring effort and capital into things that do not show up on the balance sheet yet: cleaned-up data pipelines, redesigned processes, retrained people, new governance. Those investments are real, and they are what makes later value possible, but during the building phase the metrics look flat or worse. Then, as that groundwork matures, the curve turns up, sometimes sharply, and early movers pull ahead.

Recent work extends the pattern to industrial AI specifically. Analyses of U.S. Census Bureau data published in 2025 confirm the same shape in early industrial AI adoption: short-run disruption followed by medium-term improvement for most firms.



**Figure 3.** The AI value J-curve. The dip is not failure. It is the cost of building the foundation, and it is the precondition for everything after it. Source: Brynjolfsson, Rock & Syverson (2021); McElheran et al. (2025).

**The dip is not failure.** It is the cost of building the foundation. If your Scorecard has you early in the journey, expect the first 12 to 18 months to be about laying pipe: connecting systems, cleaning data, standardizing definitions, building trust. That work will not show up as margin gains on day one. It is still the highest-value work available, because it is the precondition for everything after it.

**The plants that invest early sit higher on the curve when the value arrives.** The reason to start now is not fear of a closing window. It is compounding. Every month of clean, connected, governed data is a month your future analytics and AI can actually learn from. That history cannot be bought back later.

## What good looks like, and your next move

*Find your tier in each category and act on the single next rung. The ladder is sequential on purpose, so resist the urge to jump two levels at once.*



Whatever the product, the same seven categories determine whether a plant's data can be trusted and used.

### 1. Data Foundation & Infrastructure

This is the bedrock: where your data lives, whether it is captured reliably, whether it sits somewhere durable and accessible, and whether there is one version of the truth rather than fifteen competing spreadsheets. Everything else in the Scorecard sits on top of this, and when people describe why AI projects fail, they are usually describing a failure here.

A mature plant captures data automatically from equipment rather than by hand, keeps it in a centralized and governed store that serves as the single source of truth, and applies structure consistent enough that a batch or a downtime event means the same thing on Line 1 as on Line 4. The industry has a long way to go: one 2024 report found 35% of manufacturing processes still run on paper and 49% are tracked with manual spreadsheets, and data silos remain the top data-management concern, cited by 68% of organizations in a 2024 survey. When Nestlé rebuilt

its foundation, it retired 17 legacy siloed systems and cut the time to ingest new data sets in half.

|                         |                                                                                                                                                                                                                                                                                       |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ● <b>Reactive</b>       | Data lives on clipboards, whiteboards, and local spreadsheets. Knowledge lives in people's heads. <i>NEXT pick your two or three most decision-critical streams, starting with production output and downtime, and get them captured consistently in one shared, backed-up place.</i> |
| ● <b>Connected</b>      | Key systems capture data digitally but do not talk to each other. <i>NEXT stand up a central repository and start landing your priority sources into it. Define what a single source of truth means before you build one.</i>                                                         |
| ● <b>Centralized</b>    | A warehouse or lakehouse holds the major sources, but structure and quality are inconsistent. <i>NEXT standardize definitions and master data, products, lines, reason codes, so the data is comparable across the plant.</i>                                                         |
| ● <b>Insight-Driven</b> | Clean, governed, well-structured data supports self-service analysis. <i>NEXT extend automated capture to the last manual holdouts and harden quality gates so the foundation can support predictive work.</i>                                                                        |
| ● <b>Intelligent</b>    | Automated pipelines with built-in quality gates feed models continuously. <i>NEXT treat the foundation as a product with named owners and service levels, and hold definitions consistent as you expand across sites.</i>                                                             |

## 2. Connectivity & Integration

Connectivity is whether your systems and machines exchange data automatically, on the plant floor and between the floor and the business. Integration is what turns a pile of isolated sources into a coherent picture. Without it, people become the integration layer, stitching reports together by hand before every review.

A well-integrated plant connects its operations technology and information technology so that production reality and business systems reflect each other in near real time. Rockwell's 2025 data shows why this is hard: 21% of manufacturing leaders name integration challenges as a top internal obstacle. The mature version of this looks like Kraft Heinz's Lighthouse control tower, built with Microsoft on Azure, which pulls internal production and inventory data together with retailer point-of-sale signals and outside factors like weather into one source of truth, shifting the supply chain from reactive to proactive.


|                   |                                                                                                                                                                                               |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ● <b>Reactive</b> | Systems are islands. Integration means someone exporting a spreadsheet and re-keying it elsewhere. <i>NEXT map your data flows and automate the single highest-pain manual handoff first.</i> |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

|                         |                                                                                                                                                                                                                    |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ● <b>Connected</b>      | Some point-to-point integrations exist, but they are brittle and undocumented.<br><i>NEXT replace the most fragile hand-built links with a maintainable integration layer, and document what connects to what.</i> |
| ● <b>Centralized</b>    | Core systems are integrated through a hub, but the plant floor and the business are still separate worlds. <i>NEXT bridge the two so plant-floor data reaches business systems automatically.</i>                  |
| ● <b>Insight-Driven</b> | The floor and the business are connected, and data moves in near real time.<br><i>NEXT extend integration outward to suppliers and customers, pulling in inbound quality data and retailer signals.</i>            |
| ● <b>Intelligent</b>    | End-to-end connectivity from supplier to shelf feeds a control-tower view.<br><i>NEXT add closed-loop automation, where integrated data triggers action rather than just visibility.</i>                           |

### 3. Visibility & Analytics

Visibility is whether people can see what is happening in the plant, ideally in real time. Analytics is whether they can understand why. This is where OEE lives, and OEE remains the single most useful lens most plants have for turning data into money.

World-class OEE is widely cited at 85%, made up of roughly 90% availability, 95% performance, and 99.9% quality. Most food and beverage plants run between 55% and 65%. That gap is the opportunity, and it is a large one: moving a packaging line from 60% to 75% OEE adds the equivalent of 15% more capacity with no capital spend. Two things worth knowing before you lean on that 85% figure. It came from discrete automotive manufacturing; a realistic world-class ceiling for food and beverage sits closer to 82%, because F&B carries mandatory sanitation, allergen changeover, and quality-hold time that discrete plants do not. And manually reported OEE tends to run 8 to 12 points higher than sensor-measured OEE, so a plant that believes it is at 78% may actually be at 58%. The first honest step in visibility is measuring OEE from the machines, not from a spreadsheet.



**Figure 4.** Most F&B plants run 55 to 65% OEE against a realistic ceiling near 82%. Every point on a bottleneck line has a dollar value attached to it.

|                         |                                                                                                                                                                                                                                           |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ● <b>Reactive</b>       | Performance is known after the fact, at week's end or month's end. <i>NEXT establish an honest, sensor-based OEE baseline on your bottleneck line. Expect it to land lower than the spreadsheet says. That gap is your business case.</i> |
| ● <b>Connected</b>      | Dashboards exist but are backward-looking and often disputed on the floor. <i>NEXT put real-time OEE on a screen operators can see during the shift, so losses get caught while they are still fixable.</i>                               |
| ● <b>Centralized</b>    | Real-time visibility exists across lines from a single source. <i>NEXT move from descriptive to diagnostic analytics, starting with automated downtime reason-code analysis.</i>                                                          |
| ● <b>Insight-Driven</b> | Analytics routinely explain root causes and put a dollar figure on losses. <i>NEXT push analytics upstream toward prediction, yield, giveaway, quality drift, on your highest-value line.</i>                                             |
| ● <b>Intelligent</b>    | Analytics are embedded in the flow of work and increasingly predictive. <i>NEXT close the loop so insights trigger scheduling, maintenance, and process changes on their own.</i>                                                         |

## 4. Advanced Analytics & AI

This is predictive and prescriptive work: forecasting demand, predicting equipment failures, optimizing schedules, catching quality defects with computer vision, using generative AI to speed up knowledge work. It sits fourth in this sequence because it depends the most on the three categories before it.

McKinsey reports that AI can cut demand-forecast error by 20 to 50%. Nestlé cut its own forecast error by 30% and now reports 94% forecast accuracy. On predictive maintenance,

McKinsey's documented range is 30 to 50% less unplanned downtime and 18 to 25% lower maintenance cost, with proactive repairs running four to five times cheaper than emergency ones. IoT Analytics found that 95% of predictive-maintenance adopters report positive ROI, 27% of them within twelve months. PepsiCo's Frito-Lay, working with a machine-health monitoring partner across 38 plants, avoided more than 4,500 hours of unplanned downtime across more than 900 events.

**One figure to leave out of your own case.** *The widely repeated claim of \$26 million from predictive maintenance is an illustrative calculation, roughly 100 downtime hours at \$260,000 an hour, not a documented result at a named company. Use the ranges above and the PepsiCo case instead. It is exactly the kind of number a hype deck quotes and a credible operator should not.*

|                         |                                                                                                                                                                                                                                                                           |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ● <b>Reactive</b>       | No predictive work; analytics, where they exist, only look backward. <b>NEXT</b> <i>do not start here. Fix categories one through three first. The most valuable thing you can do for a future AI project is give it clean, connected data to learn from.</i>             |
| ● <b>Connected</b>      | Curiosity exists, and maybe a proof of concept ran and quietly stalled. <b>NEXT</b> <i>run one tightly scoped pilot with a defined success metric and a data-readiness check up front, on a problem where your data is already good, usually downtime or forecasting.</i> |
| ● <b>Centralized</b>    | One or two models are in production and trusted by the people using them. <b>NEXT</b> <i>standardize how a model moves from pilot to production, and document the data and monitoring each one needs to stay reliable.</i>                                                |
| ● <b>Insight-Driven</b> | Multiple models run in production across different functions. <b>NEXT</b> <i>build a repeatable practice for maintaining models so they get monitored instead of left to drift.</i>                                                                                       |
| ● <b>Intelligent</b>    | AI, including generative and causal AI, is a core operating capability. <b>NEXT</b> <i>pursue prescriptive and increasingly autonomous decisions where the foundation is strong enough to support them safely.</i>                                                        |



Automation is only as good as the data feeding it. The nozzle is the easy part.

## 5. Governance, Data Quality & Security

Governance is who owns the data, how quality is maintained, and how it is secured. It is the least visible category here and one of the most decisive. Informatica's 2025 survey put data quality and readiness at the very top of the obstacles to AI success. Governance is not paperwork for its own sake. It is what makes data trustworthy enough to act on, and what keeps a newly connected plant from becoming an easy target.

A governed plant has named owners for its key data domains, defined quality standards with automated checks behind them, master data management so products and locations and reason codes stay consistent, and security appropriate to a plant where operations technology and information technology are now linked. That last point gets more important as you climb the ladder, since connectivity widens the attack surface. Rockwell's 2025 data ranks cybersecurity as the second biggest external risk manufacturers see, and 49% plan to use AI for cybersecurity itself in 2025. McKinsey's research backs up why this matters: organizations reporting real financial returns from AI are twice as likely to have redesigned their data workflows before they ever chose a modeling technique.

|                         |                                                                                                                                                                                                                                     |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ● <b>Reactive</b>       | No formal ownership; quality is whatever the last person entered; security is thin. <i>NEXT assign clear ownership for your two or three most important data domains and write down the quality rules that actually matter.</i>     |
| ● <b>Connected</b>      | Standards exist on paper but are not enforced day to day. <i>NEXT add automated quality checks at the point of capture and close the security gaps your new connectivity opened up.</i>                                             |
| ● <b>Centralized</b>    | Governance exists for core domains and is actually enforced. <i>NEXT implement master data management so definitions hold across systems and sites.</i>                                                                             |
| ● <b>Insight-Driven</b> | Governance is systematic and mostly automated. <i>NEXT extend governance to your models too: versioning, monitoring, and quality checks for the analytics you now depend on.</i>                                                    |
| ● <b>Intelligent</b>    | Governance runs continuously across the enterprise, covering data and models alike. <i>NEXT push toward data that is aligned to specific use cases and governed at the asset level, with continuous quality assurance built in.</i> |

## 6. People, Skills & Strategy

Technology does not transform a plant. People using it do. This category covers whether you have the skills, the change-management muscle, and the strategy to turn tools into outcomes, and it is where the J-curve is actually lived, since the groundwork that curve describes is mostly investment in people and process.

Rockwell's 2025 data shows why the mature organizations treat this as first-class work: manufacturers named resistance to change and leading change as their two biggest leadership obstacles, and nearly half of respondents now call the ability to apply AI an extremely important skill, up from just 10% a year earlier. In food and beverage specifically, producers ranked training current employees, retention, and knowledge capture as their top three workforce concerns for the year ahead. That matters acutely in an industry losing institutional knowledge as experienced operators retire. Nestlé's own framing is useful here: it treats AI as something that enables people, not something that replaces the need to train them, and anchors every use case in an actual business need rather than the technology itself.

|                         |                                                                                                                                                                                                                          |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ● <b>Reactive</b>       | Data and analytics are one person's side project; no real strategy exists. <i>NEXT get executive sponsorship, name a single owner for the data journey, and write a one-page strategy tied to actual business goals.</i> |
| ● <b>Connected</b>      | Pockets of skill exist but nobody is coordinating them. <i>NEXT build basic data literacy across supervisors and start capturing tribal knowledge before it walks out the door.</i>                                      |
| ● <b>Centralized</b>    | A defined team and roadmap exist. <i>NEXT invest in change management as deliberately as you invest in tools. The technology is rarely what actually stalls a project.</i>                                               |
| ● <b>Insight-Driven</b> | Data-informed decisions are simply how the organization works. <i>NEXT develop advanced analytical and AI skills in-house and build career tracks that keep that talent from leaving.</i>                                |
| ● <b>Intelligent</b>    | A data-and-AI culture runs from the floor to the leadership team. <i>NEXT sustain it through continuous learning, cross-training, and knowledge systems that outlast any one person's tenure.</i>                        |



Every can, bottle, or case that leaves the plant carries a lot code. Traceability is whether that code actually means something.

## 7. Traceability & Compliance

Traceability is the ability to follow product and its inputs through your operation and produce the records fast when it matters. For F&B manufacturers, this ties directly to recall exposure and to FSMA 204. The financial logic is stark: the commonly cited average direct cost of a food recall is about \$10 million, covering only direct costs like retrieval and disposal. In the underlying industry study, 52% of companies that experienced a major recall reported total impact above \$10 million, and 1 in 20 exceeded \$100 million, before counting lost contracts or brand damage.

The FDA's Food Traceability Final Rule requires companies that manufacture, process, pack, or hold foods on the Food Traceability List to keep records of Key Data Elements tied to Critical Tracking Events, and to hand them to the FDA in an electronic, sortable format within 24 hours of a request. Critical Tracking Events are the points where food is handled, transformed, or moved: harvesting, cooling, initial packing, shipping, receiving, transformation. Key Data Elements are the specifics attached to each one: traceability lot codes, product descriptions, quantities, dates, locations. The compliance date was originally January 20, 2026, but the FDA pushed it back in 2025, and the deadline now sits at July 20, 2028. That extension is time to build, not a reason to stop. The requirements have not changed, and many retailers already expect traceability turnarounds faster than the FDA's 24 hours.

*FSMA 204 is, functionally, a data-management mandate. It requires clean, standardized, linkable data that can be produced fast. That is the same foundation AI needs.*

Build traceability well, and you build much of the substrate for analytics at the same time.

- |                         |                                                                                                                                                                                                                     |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ● <b>Reactive</b>       | Traceability is manual, and a mock recall would take days. <i>NEXT confirm which products are on the Food Traceability List, map your Critical Tracking Events, and assign traceability lot codes consistently.</i> |
| ● <b>Connected</b>      | Records are digital but scattered, so assembling them takes time. <i>NEXT centralize traceability records so Key Data Elements link across events and can be pulled quickly.</i>                                    |
| ● <b>Centralized</b>    | You can produce records within the required window, with real effort. <i>NEXT automate KDE capture at each event and run a timed mock traceback to find the gaps.</i>                                               |
| ● <b>Insight-Driven</b> | Traceability is largely automated and fast. <i>NEXT extend the digital thread to trading partners so inbound and outbound data flows cleanly.</i>                                                                   |
| ● <b>Intelligent</b>    | End-to-end traceability runs in real time and works across partners. <i>NEXT use the traceability data itself as an analytics asset for quality and yield, not just compliance.</i>                                 |

# Documented results, and the cost of standing still

*Figures you can take to a CFO, with the source and the caveat attached where one is warranted.*

## Real, documented outcomes


| COMPANY             | WHAT HAPPENED                                                                                                                                                                                                                                                                |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Nestlé              | Digital adoption returned about 1.5 million productivity hours, reported at roughly \$30 million. Machine learning cut demand-forecast error by 30%, with accuracy now reported at 94%. A centralized data lake eliminated 17 legacy systems and cut ingestion time in half. |
| Kraft Heinz         | Its Lighthouse control tower, built with Microsoft on Azure, aggregates internal, retailer, and external data into one source of truth, shifting the supply chain from reactive to proactive. A foundation-and-integration story more than an AI-first one.                  |
| General Mills       | AI applied to logistics assessed more than 5,000 daily plant-to-warehouse shipments and drove over \$20 million in savings since fiscal 2024, with real-time manufacturing data expected to add more than \$50 million in waste reduction.                                   |
| PepsiCo / Frito-Lay | Machine-health monitoring across 38 plants avoided more than 4,500 hours of unplanned downtime across 900-plus events. Quantified in recovered hours and capacity, not a single dollar figure.                                                                               |

Sources: Nestlé 2025 annual report and Fortune, 2026; CIO Dive, 2025; Forbes, 2024. A McKinsey estimate puts the potential value of a full digital and AI transformation for a \$10 billion food and beverage company between \$810 million and \$1.6 billion. That figure is an estimate of potential, not a realized result, and is labeled as such here on purpose.

|                                                                                                        |                                                                                    |                                                                                                                    |
|--------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| <b>20 to 50%</b><br>forecast-error reduction from AI, flowing straight into working capital (McKinsey) | <b>30 to 50%</b><br>less unplanned downtime from predictive maintenance (McKinsey) | <b>about 10%</b><br>output increase from SKU rationalization, alongside a fifth less manufacturing cost (McKinsey) |
|--------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|

## The cost of inaction

This part is worth doing carefully. The goal is not a frightening number. It is to show that inaction has a running cost, with ranges a CFO will actually respect. Replace the assumptions below with your own figures before you use them anywhere.



**Figure 5.** Illustrative relative magnitudes. Unplanned downtime is usually the largest line for a mid-size plant, but replace every input here with your own numbers before making a decision.

**Manual data labor.** When data lives in silos, people spend their days assembling it instead of using it. Independent studies from Forrester and IDC both find knowledge workers spend around 30% of their time, about 2.5 hours a day, simply looking for information. For a plant with a handful of salaried people doing this, that is a six-figure annual cost in wages alone, before counting the decisions delayed or made wrong on stale numbers.

**Unplanned downtime.** Usually the largest line. ABB's 2023 reliability survey put the median cost of downtime at \$125,000 an hour, with more than two-thirds of businesses experiencing unplanned downtime at least monthly. You do not need a headline number to make this case. Use your own line's finished-goods value per hour: a line producing \$40,000 an hour that suffers a six-hour unplanned stop just lost \$240,000 before the repair bill arrives. If predictive approaches cut unplanned downtime by even 30%, the low end of the documented range, the recoverable value is easy to calculate and easy to defend.

**The OEE gap and forecast error.** If your honest, sensor-measured OEE sits at 60% against a realistic ceiling near 80%, that gap is sellable product you are not making on equipment you already own. And every point of forecast error shows up twice, once as excess inventory tying up working capital, and again as stockouts costing sales.

**Recall exposure.** Tie this to readiness rather than fear. The average direct recall cost is about \$10 million, and nearly a quarter of recalls in the industry study exceeded \$30 million, before counting lost contracts or brand damage. The variable you actually control is speed. A plant that can trace a lot in minutes contains a recall. A plant that needs days recalls far more than it should.

## Sequencing, and your first moves

*The categories are not a menu. They are a sequence.*

Build in this order. First, **data foundation**: get critical data captured, stored, and trusted in one place. Nothing else works until this does. Second, **connectivity and integration**: make the systems talk, then bridge the plant floor and the business so they reflect each other. Third, **visibility and analytics**: put honest, real-time performance in front of people, starting with sensor-measured OEE on the bottleneck. Fourth, and only now, **advanced analytics and AI**: layer in prediction and optimization on the problems where the foundation is strongest.

Two categories run alongside all four. **Governance, data quality, and security** has to mature in step with each rung, since every added level of connectivity and analytics raises the stakes on trust and security. **People, skills, and strategy** is the engine underneath the whole climb, and it is where the J-curve's early investment actually happens. **Traceability and compliance** is both an obligation and something of a gift: building it well builds much of the foundation the other categories need, and the 2028 FSMA 204 deadline gives you room to treat it as infrastructure rather than a scramble.

Here is how to read your Scorecard against this sequence. Find your lowest-scoring foundational category, one of the first three. That is almost certainly where your next dollar and next month belong, whatever a flashier project higher up the list might promise. If category four, AI, is your weakest score but the first three are also weak, the right move is counterintuitive but well supported by the evidence: do not start the AI project yet. Fix the foundation first, and the AI project you run in twelve months will be the one that survives contact with production.

### A staged set of first moves

#### NEXT 90 DAYS · SEE THE TRUTH

- Establish an honest, sensor-based OEE baseline on your single most important line. Expect it to land 8 to 12 points below the spreadsheet. That gap is your business case.
- Identify your lowest-scoring foundational Scorecard category and scope one concrete project to move it up a single level.
- Confirm which products sit on the Food Traceability List, map your Critical Tracking Events, and run a timed mock traceback. If it takes more than 24 hours, traceability moves to the top of your list.

### 3 TO 12 MONTHS · BUILD THE FOUNDATION

- Stand up or consolidate a single source of truth for your priority data streams, and standardize master data: products, lines, downtime reason codes.
- Automate your highest-pain manual data handoff, and bridge one gap between the plant floor and the business so data flows automatically.
- Assign ownership and quality rules for your two or three most important data domains, and close the security gaps your new connectivity opened.

### 12 MONTHS AND BEYOND · EARN YOUR WAY TO AI

- Run one tightly scoped predictive pilot, with a defined success metric and a data-readiness check up front, on your best-instrumented problem.
- Build a repeatable path from pilot to production so you do not join the majority that stall, and invest in change management and knowledge capture alongside it.
- If a pilot cannot show measurable value inside its defined window, stop and check the data foundation before blaming the model. That is usually where the fault sits.

*Haystack's approach is to build with you, not for you: diagnose first, then prescribe, and leave your team more capable than we found it. The Scorecard is the diagnose step.*

If it would help to turn your result into a costed, staged roadmap, that is what our AI Readiness work is for. No urgency, no closing window. Just the next right rung, when you are ready for it.

### A few honest caveats

- **Market-size forecasts vary widely and should be read as directional only.** Published estimates for the AI-in-F&B market range from roughly \$10 billion to figures implying \$84 billion or more by 2030, with growth-rate estimates from about 15% to over 40%. That spread reflects differing definitions, not precision, which is why this report leans on operational and adoption data instead.
- **Some figures come from vendors or secondary sources.** The Nestlé, General Mills, and McKinsey figures are the best sourced here. Treat single-source vendor case studies with more caution than named-company annual reports.
- **AI failure-rate statistics carry a range.** The 80% figure from RAND and the 95% figure from MIT measure different things, projects that never reach production versus pilots that show zero measurable return. Both point the same direction without being a single precise truth.
- **OEE benchmarks depend on measurement method and subsector.** The 85% figure has its roots in discrete manufacturing. Realistic F&B ceilings run lower, and manual reporting inflates scores, so any OEE comparison is only as good as the measurement behind it.

## Sources

---

*Everything here is dated 2023 or later. Primary and reputable sources were preferred throughout.*

- Rockwell Automation.** 2025 State of Smart Manufacturing Report, and the Food & Beverage / CPG editions; Modernization in Food and Beverage Manufacturing (2025 to 2026). [rockwellautomation.com](https://rockwellautomation.com)
- RAND Corporation.** The Root Causes of Failure for AI Projects (RR-A2680-1), Ryseff, De Bruhl & Newberry, 2024. [rand.org](https://rand.org)
- McKinsey & Company.** Fortune or fiction? The real value of a digital and AI transformation in CPG (2024); What it takes to rewire a CPG company (2024); State of Food and Beverage (2025 to 2026). [mckinsey.com](https://mckinsey.com)
- MIT Project NANDA.** The GenAI Divide: State of AI in Business 2025.
- S&P Global Market Intelligence.** Voice of the Enterprise: AI & Machine Learning, 2025.
- Informatica.** 2025 CDO Insights survey.
- Brynjolfsson, Rock & Syverson.** The Productivity J-Curve: How Intangibles Complement General Purpose Technologies, NBER/AEA, 2021. McElheran et al., industrial AI adoption using U.S. Census microdata, 2025.
- U.S. FDA.** FSMA Final Rule on Requirements for Additional Traceability Records for Certain Foods; Federal Register compliance-date extension, 2025. [fda.gov](https://fda.gov), [federalregister.gov](https://federalregister.gov)
- Food Safety Magazine.** FDA Delays FSMA 204 Traceability Rule Compliance Date by 30 Months, 2025.
- Nestlé.** Accelerating our business transformation, 2025 annual report; Nestlé USA, Unlocking New Opportunities with Gen AI; Fortune, Nestlé's CIO on the value of AI investments, 2026; Emerj, Artificial Intelligence at Nestlé.
- Food Dive / Supply Chain Dive.** AI, automation enrich Kraft Heinz's supply chain, 2024. Forbes, How Kraft Heinz Is Using Artificial Intelligence, 2024.
- CIO Dive / Food Dive.** General Mills attributes millions in cost savings to AI, 2025.
- Forbes.** How PepsiCo Avoids Pilot Purgatory With Innovation Partnerships, 2024.
- Predictive maintenance and downtime.** McKinsey benchmarks; ABB Value of Reliability Survey, 2023; Siemens True Cost of Downtime, 2024; IoT Analytics, 2023.
- OEE benchmarks.** Nakajima TPM; TeepTrak, iFactory, OxMaint, Traction, 2026 syntheses.
- Recall costs.** Grocery Manufacturers Association / Covington & Burling / Ernst & Young food-recall cost study; GMA/FMI average recall-cost figures.
- Data silos and manual process prevalence.** DATAVERSITY, 2024 Trends in Data Management; iBASeT/MTC Digital Manufacturing Productivity Report; L2L, 74% of Manufacturers Caught in Data Chaos, 2025.
- Knowledge-worker time.** Forrester, commissioned by Airtable, and IDC studies on time spent searching for information.
- Beyond pilots.** Consumer Goods Technology, Where Food & Beverage Manufacturers See Real ROI from AI, 2025.

---

This report is vendor-neutral and educational, built for leaders who completed the F&B Plant Data Maturity Scorecard. Figures are drawn from sources dated 2023 or later. Replace illustrative assumptions with your own plant's numbers before making an investment decision. Prepared by Haystack Data Solutions, a food intelligence company. © 2026 Haystack, LLC.