

Honest AI

April 2026 | Edition 15

From Gut to Algorithm

How AI is changing the game for CPG brands and food processors

Webinar on:

How to pass your food safety audit with AI – without the panic

Most CPG brands and food processors spend days scrambling to prepare for an audit. What if AI could cut that down to minutes? Join us and find out exactly how.

Nishkam Batta, Editor-in-Chief, HonestAI Magazine · As seen in Forbes, Entrepreneur, Morning Brew

May 6th, 2026 · 11AM EST



Register for webinar

Table of Contents



The founder algorithm: Why your gut got you here but AI will get you there

1

- 1.1 Founder instinct is a system but it lives in one brain
- 1.2 AI doesn't replace founder instinct, it encodes It



Lot tracking, reinvented: How AI is solving the food industry's oldest data problem

12



The Co-Packer intelligence gap: How AI bridges the data divide without asking anyone to change

18


- 3.1 Validation rules that matter most
- 3.2 Meeting Co-Packers where they are



When your spreadsheet thinks: How AI is transforming the master order sheet into a living system

26

- 4.1 What spreadsheets do better than most systems
- 4.2 When recalls become crisis: The role of data precision in containing risk



AI on the floor: What's working, what's not, and what every food & beverage operator needs to know right now

45

- 5.1 Where AI is actually delivering ROI today
- 5.2 The role of AI middleware



Winning back the deduction: How AI-powered documentation is turning distributor disputes in favor of small brands

56

- 6.1 How distributor deductions actually work
- 6.2 The financial impact of deduction recovery



Scaling the kitchen: How AI helps scale beyond your first Co-Packer

62



The blueprint before the build: Why the most important AI investment you'll make is documentation

67



LETTER *from* The Editor In Chief

Dear Readers,

The food and beverage industry has always been built on instinct.

Founders trust their gut to develop products, spot trends, and navigate early growth. In the beginning, that instinct is often enough. It drives decisions, solves problems, and holds the entire operation together.

But as brands scale, something changes.

What once felt manageable becomes fragmented. Orders live in one system. Inventory in another. Co-packer's production data in spreadsheets. Retail demands shift faster than the systems tracking them. When USDA audits knock on the door, the entire company is preparing for it for weeks.

The business grows, but so does the complexity behind it. At that point, instinct alone starts to reach its limits. This issue explores a shift that is already underway.

AI in food and beverage is not a future concept. It is already embedded inside the operations of brands that are quietly pulling ahead, helping them connect systems, validate data, and make decisions faster than manual processes ever could.

What's important is this: adopting AI does not

mean replacing everything you have built. It means connecting it.

The most effective companies are not starting from scratch. They are taking existing spreadsheets, order records, HACCP tools, ERPs, Co-Packer logs and layering intelligence on top, turning disconnected tools into coordinated systems.

That is the transition from gut to algorithm.


From reacting to problems > to preventing them.
From chasing data > to trusting it.
From managing operations > to designing them.

This issue is for the founder who turned a recipe into a business and is now realizing the business has outgrown the way it's being run.

There is a better way. And it starts with understanding where you stand today. That's what the AI Readiness Assessment is built to answer.

Take our 5-minute AI Readiness Assessment to see how prepared your operation is and where the biggest opportunities lie. It is a diagnostic assessment that tells food brands and processors exactly where they stand on AI readiness and what it's costing them not to act. It produces a maturity stage, a projected ROI exposure range, and a clear picture of where operational friction is highest.

Sincerely,
Nishkam Batta
Editor-in-chief, HonestAI Magazine





01

The founder algorithm: Why your gut got you here but AI will get you there

Every successful food and beverage founder runs an algorithm. Not on a laptop. Not in an ERP system.

In their head, they know which lot to pull before QA even raises a concern. They know which Co-Packer can deliver under pressure and which one will quietly miss timelines. They know which distributor deductions are legitimate and which ones deserve a dispute.

This instinct isn't magic. It's a pattern recognition

built through years of operational experience.

But here's the reality most founders eventually face: what works when the company is small becomes the biggest bottleneck when the company grows.

The operational intelligence that once made the founder indispensable becomes the very thing that limits scale.

The brands winning today are not replacing founder instinct with AI. They're encoding it.

1.1

Founder instinct is a system but it lives in one brain

Early-stage food and beverage companies rarely begin with perfect systems.

They begin with people. Especially founders.

A founder quickly becomes the operational center

of gravity for the entire business. Information flows through them because they're the only person who understands how everything connects. Over time, they develop a powerful mental model of the business. It works like an invisible operating system. For example, a founder might instinctively know:



Which production runs are risky

A founder may notice that a certain ingredient supplier often causes inconsistencies in texture or shelf life. Even if lab results look fine, their experience tells them the batch needs extra attention.



Which distributors require scrutiny

In CPG, distributor deductions can quietly erode margins. Experienced founders know which deductions are legitimate and which ones are routine errors that must be challenged.



Which retailers reorder reliably

Founders often understand which retail partners generate consistent sell-through and which ones place large opening orders but struggle to move product.



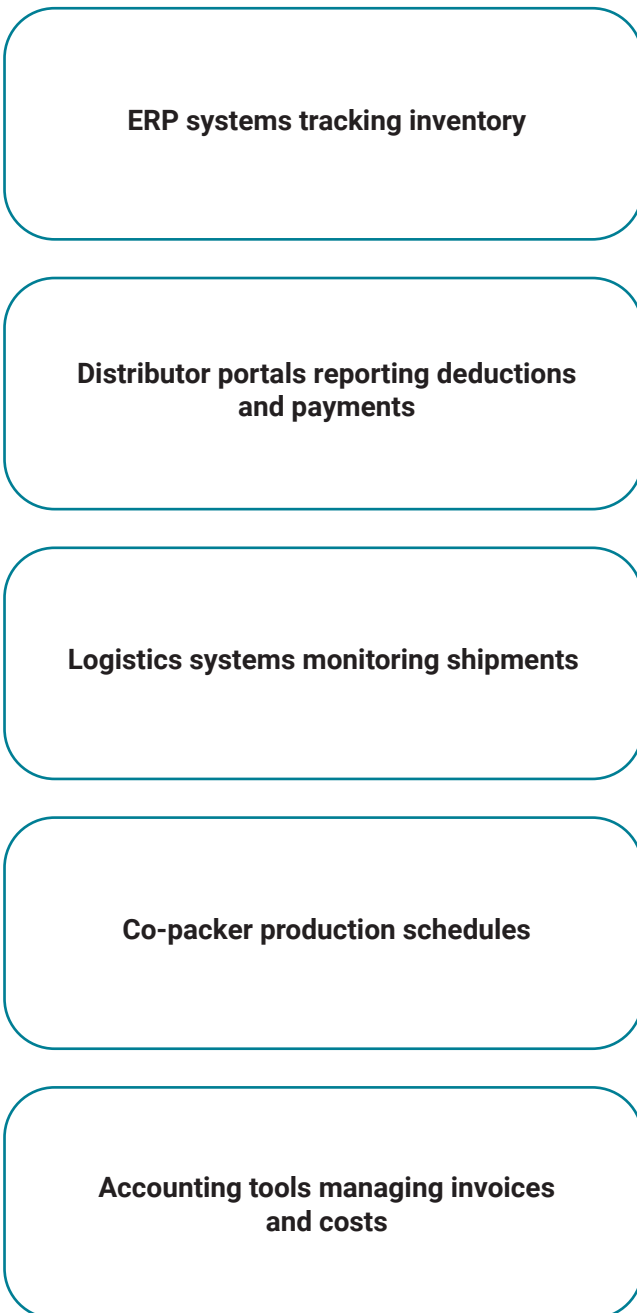
When inventory signals don't make sense

Inventory numbers may look correct in software, but founders often detect inconsistencies based on shipment timing, production yields, or historical patterns.

These are not random decisions. They are pattern recognition and exception handling, the same core logic used by modern AI systems. The difference is that founders do it manually.

The problem: Human intelligence doesn't scale easily

As the company grows, complexity grows faster than instinct can keep up. Suddenly the founder is juggling information from dozens of sources:



Each platform holds a piece of the operational puzzle. But none of them see the whole picture. So the founder becomes an integration layer.

They are the one connecting signals between systems and resolving conflicts when something doesn't line up.

For example:

A distributor reports that inventory was short shipped.

The ERP shows the order was fulfilled correctly.

The Co-Packer production report suggests the yield may have been lower than expected.

Three systems. Three different versions of reality.

The founder becomes the detective who figures out what actually happened. This kind of operational detective work is common across growing CPG brands.



Scan to watch: He sold nut butter for \$281M

According to APQC benchmarking, lower-performing order-to-cash operations may require manual intervention on 20% or more of orders, roughly one in five. In founder-led companies, that burden often falls on the founder in the early stages.

Justin Gold's early journey with Justin's reflects the kind of hands-on founder role common in emerging food brands. Gold built the company in Boulder, sold products at farmers' markets, secured local retail placement including Whole Foods, and continued working other jobs while the business was still developing.

In 2016, Hormel acquired Justin's for a reported purchase price of \$280.9 million.

Justin's illustrates a familiar CPG pattern: early growth often depends on unusually high founder involvement, but scale eventually requires more formal operating systems, clearer workflows, and less dependence on one person's direct oversight.

This final sentence is an interpretation drawn from the broader pattern, not a direct claim about internal Justin's systems.

1.2

AI doesn't replace founder instinct, it encodes it

For many founders, the first reaction to AI is Uncertainty.

Food and beverage businesses are built on judgment that comes from years of navigating messy supply chains, unpredictable retailers, and production realities that rarely match what spreadsheets promise. Founders develop instincts about their operations that no dashboard can fully capture. They know when something “doesn't feel right” long before the numbers confirm it.

That's why the idea of AI replacing founder decision-making often feels unrealistic. But the real value of AI isn't replacing instinct. It's capturing it and making it scalable.

Every operational decision a founder makes, whether it's identifying a risky production batch, challenging a questionable distributor deduction, or deciding which order should move first. These

decisions aren't random.

They're shaped by years of experience, repeated exposure to similar situations, and lessons learned from past mistakes.

AI systems operate in much the same way. They analyze historical data, recognize patterns across operations, and flag anomalies when something doesn't align with expected behavior. The difference is scale. While a founder can track dozens of signals at once, AI can monitor thousands of operational variables simultaneously and do it continuously.

For growing CPG brands, this shift is essential. The goal isn't to remove the founder from the operational loop. It's to ensure the business can operate with the same intelligence, judgment, and responsiveness even when the company grows to ten times its current scale.

What AI actually does inside food operations

When founders hear “AI,” they often imagine complex prediction models or futuristic automation. But inside most food and beverage operations, AI plays a much simpler role.

It acts as a continuous monitoring layer across the systems that already run the business – ERP, inventory platforms, distributor portals, retail data feeds, and logistics systems.

At its core, AI helps teams do three practical things:

1

Track lot numbers and batch data in real time

AI continuously monitors lot numbers across production, storage, and distribution. Instead of manually tracing batches, teams can instantly see where a specific lot is from raw materials to finished goods in the market.

2

Improve traceability and recall readiness

With AI connecting data across systems, traceability becomes faster and more reliable. If there's a quality issue or recall, teams can quickly identify affected batches, trace their movement, and take action without delays or guesswork.

3

Ensure compliance and visibility across the supply chain

From expiration dates and batch tracking to supplier records and distribution logs, AI helps maintain a clear audit trail. This reduces compliance risks and gives teams full visibility into how products move through the supply chain.

4

Recognize patterns in operational data

AI systems analyze historical operational data to understand what "normal" looks like for a business. For example, they may learn typical production yields from a Co-Packer, standard shipment timelines to certain retailers, or how distributor orders usually correlate with sell-through at the store level.

Once those patterns are established, the system can compare new activity against that historical baseline. If a distributor suddenly places an order that's significantly larger than typical demand patterns, or if a production batch yield drops below its normal range, the system can highlight that change for review.

The goal isn't to make the decision automatically, it's to surface signals that humans might otherwise miss.

5

Identify operational inconsistencies early

Food supply chains generate many small errors that often go unnoticed until they become bigger issues. These might include mismatched order quantities, incorrect distributor deductions, shipment delays, or inventory discrepancies between systems.

AI tools can continuously check for these inconsistencies across multiple data sources. For instance, if a distributor reports receiving fewer units than the shipment record shows, the system can flag the discrepancy immediately rather than weeks later during reconciliation.

By catching these issues early, teams can resolve them before they affect retailer relationships, inventory availability, or financial reporting.

6

Direct issues to the right people faster

In many growing CPG companies, operational issues bounce between teams before they reach the person who can actually resolve them.

AI systems can help by automatically routing flagged issues to the relevant team or system. For example, a production yield anomaly might be sent directly to the operations team, while a questionable deduction could be forwarded to finance for review.

This doesn't remove human decision-making, it simply shortens the time it takes for the right person to see the problem.

In practical terms, these capabilities allow food businesses to move from reactive problem-solving to earlier operational awareness by identifying issues while they're still small enough to fix easily.

And that's often where the real operational value of AI appears first: not in replacing decisions, but in making sure the right decisions happen sooner.

Example 1: Distributor deduction validation



Distributor deductions are one of the most common and often overlooked sources of revenue leakage in the CPG industry. Many brands only discover incorrect deductions weeks or months later during financial reconciliation, when the trail is harder to trace and the recovery window has already narrowed.

This is where an AI-powered operational

intelligence layer from GrayCyan can make a meaningful difference.

Instead of relying solely on manual reviews, GrayCyan's AI monitors operational data across the systems that food brands already use including Co-Packer logs, ERP platforms, HACCP tools, sales spreadsheets, distributor portals, trade promotion records, shipment confirmations, and accounting systems.

For example, the system can continuously analyze:

Trade promotion agreements:

Verifying that discounts or allowances match the actual promotional terms.

Distributor deduction behavior:

Comparing new deductions with historical deduction patterns from that distributor.

Shipment and delivery records:

Checking whether the quantities claimed in a deduction align with confirmed deliveries.

Promotional timelines:

Validating whether a deduction is tied to an active promotion period.

If a distributor submits a deduction that doesn't align with these operational records, such as claiming a promotional discount that wasn't scheduled or applying a chargeback inconsistent with shipment data GrayCyan's AI flags the discrepancy immediately.

Instead of discovering the issue months later during reconciliation, the finance or operations team receives a real-time alert while the transaction is still recent and easier to investigate.

The goal isn't to replace the finance team's judgment. It's to surface inconsistencies early, so teams can review them quickly and take action before small issues quietly erode margins.

For many growing CPG brands operating across multiple distributors and retail partners, this kind of early validation can recover revenue that would otherwise go unnoticed while reducing the operational burden on founders who often end up investigating these issues manually.

Example 2: Inventory and production monitoring

AI systems continuously compare real-time production output with historical yield patterns and expected benchmarks. This allows them to detect even small deviations that might otherwise go unnoticed in day-to-day operations.


For example, if a batch yield suddenly drops by 8–10%, the system flags it immediately, not days later. It can also correlate that drop with specific factors like raw material lot variations, machine performance, or process inconsistencies.

That early signal gives the operations team time to act. They can adjust production schedules, inspect

specific batches, recalibrate equipment, or even redirect existing inventory to avoid downstream shortages.

In some cases, AI can also predict the impact estimating how that yield drop will affect available stock across warehouses or retail locations in the coming days.

Without that visibility, the issue often surfaces too late, typically when distributors or retailers report empty shelves or delayed orders. By then, the response becomes reactive instead of proactive, leading to lost sales and strained relationships.



A real founder scenario: RXBAR's early operational bottleneck



Scan to watch: Peter Rahal on the hard work of building RXBar | Chasing excellence

In the early days of RXBAR, Peter Rahal was not just building a brand. He and co-founder Jared Smith were deeply involved in making, selling, and growing the business themselves. Chicago Magazine reports that Rahal and Smith began making bars in Rahal's parents' Glen Ellyn basement and that after six months they moved production to a small product development facility on Western Avenue.

RXBAR's early go-to-market strategy was highly hands-on. The founders focused on direct sales, especially into CrossFit gyms, and later leaned heavily on Amazon.

Chicago Magazine reports that for the first two years, RXBAR's business plan was essentially to make a batch of product, sell it on Amazon, get paid, and then make another batch. The same report says that by managing their own inventory, the founders also "managed their own fate," suggesting an early operating model that relied heavily on direct founder involvement rather than formal systems.

Rahal's personal cell phone number appeared on early packaging, and the original wrapper was designed by the founders themselves in PowerPoint.

Those details suggest how closely the founders were tied to day-to-day execution in RXBAR's early phase.

RXBAR entered Whole Foods in 2015, followed by Trader Joe's, Kroger, and Target, with its packaging redesign helping attract major retail partners.

Chicago Magazine also reports that mass distribution increased the company's cash-flow needs, including a \$500,000 loan from Rahal's father that was later repaid. In October 2017, Kellogg announced it would acquire RXBAR for \$600 million, with expected 2017 net sales of

approximately \$120 million, and said RXBAR would continue operating independently as a standalone business.

RXBAR's early growth relied heavily on direct founder involvement across production, sales, customer contact, and inventory management. As the brand expanded into major retail channels, the scale of the business increased well beyond its original hands-on operating model. That broader lesson is supported by the reporting, even though some of the more detailed operational claims in the original draft are not.

The founder's new role

When operational intelligence becomes embedded in systems, something important changes. The founder stops running every operational decision. Instead, they begin designing the operation itself. Rather than reacting to daily operational noise, founders can focus on the areas that truly drive growth:



This shift has already happened in other industries. When cloud infrastructure automated server management in technology startups, founders stopped worrying about hardware and began focusing on building better products.

Food and beverage brands are now beginning to experience a similar transition, moving from founder-driven operations to systems that can support growth at scale.


A question every founder should ask

? What happens if the founder disappears tomorrow?


? How much operational intelligence disappears with them?

For many founder-led brands, the answer is: **a lot**


Years of experience are:



Not documented



Not systemized




Stored only in the founder's mind

This creates a hidden risk:



Business dependency on one person



Slower scaling



Operational bottlenecks

The AI opportunity


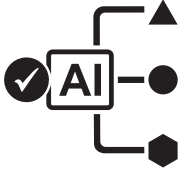
AI enables you to:



The real advantage isn't just data. It's turning instinct into infrastructure.

Key insight

The fastest-scaling brands:

- 
Don't just collect data
- 
They systemize experience



GrayCyan has built an Operational Intelligence Audit for food and beverage founders to map where operational intelligence lives inside your business today. And where AI could begin capturing it. Because your instinct built the brand. But systems are what will scale it.



02

Lot tracking, reinvented: How AI is solving the food industry's oldest data problem

Lot tracking has been a mandatory requirement in the food industry for decades. Regulations such as the U.S. Food Safety Modernization Act (FSMA) and the European Union's General Food Law require food companies to maintain traceability records that identify where ingredients come from and where finished products go.

In theory, the system is straightforward. In practice, it has long been one of the most persistent operational problems in food manufacturing.

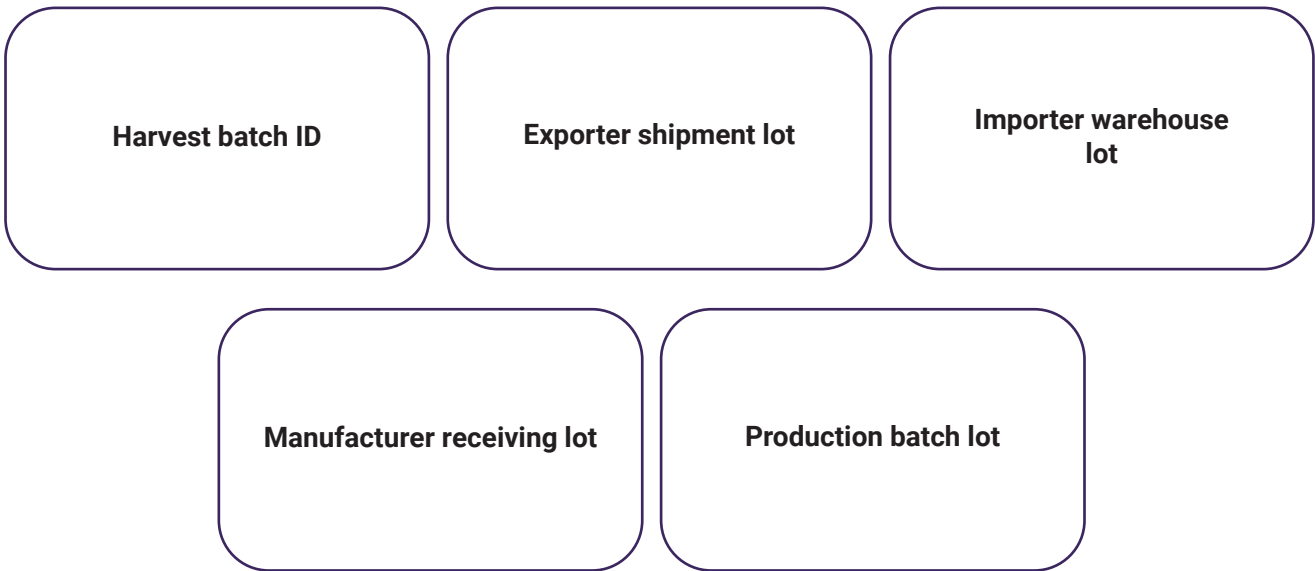
Lot numbers are created by different actors across the supply chain. Farms, processors, distributors,

warehouses, manufacturers, and Co-Packers all generate their own identifiers for the same ingredient. By the time a material reaches a production line, it may carry several different lot references.

The data is technically correct but operationally fragmented. This fragmentation leads to inventory mismatches, delayed recall investigations, and hours of manual reconciliation work. AI systems are now emerging to solve this problem by validating lot data the moment it is entered instead of discovering errors later.



A single ingredient may pass through multiple organizations before reaching a production facility. For example, a shipment of cocoa powder might carry:



Each identifier reflects a legitimate step in the supply chain. However, enterprise systems such as SAP, Oracle NetSuite, and Microsoft Dynamics typically don't store multiple identifiers by default.

The result is a familiar scenario in food plants: the inventory system shows little or no stock available even though the ingredient is physically present on the production floor.

When operators enter a different identifier during receiving or production, the system often cannot reconcile the data.

These discrepancies consume time, slow production scheduling, and complicate traceability.

Case study: ReposiTrak and automated compliance



Scan to watch: Randy fields talks about great advancements in food traceability and how the industry can comply


Another real example comes from ReposiTrak, a food traceability and compliance network led by CEO Randall K. Fields.

ReposiTrak reported in October 2024 that its Traceability Network had reached 4,000 suppliers. The company also said that supplier onboarding is entirely automated and that suppliers can exchange FDA-required traceability data with retailer and wholesaler customers through the network.

In one retailer rollout announced in June 2024, ReposiTrak said the network included about 600 suppliers and nearly 950 supplier facilities connected to transmit FDA-required Key Data Elements.

ReposiTrak has also stated that, as it implemented traceability for suppliers, up to 35% of files initially received contained some kind of error, usually missing or incorrect data, and that the company helps suppliers create cleaner files that meet customer requirements.

In a later company statement, ReposiTrak said every traceability data file is checked using a 500+ point error detection process so the data is as complete and accurate as possible before it reaches retail, wholesale, or foodservice customers.



Why catching errors early matters

Traditional systems typically discover traceability errors only after transactions fail. A simple data entry mistake during receiving such as a typo in a lot number or a misplaced digit can remain unnoticed for days or even weeks.

When production later attempts to use that ingredient, the system may report insufficient inventory or an invalid lot reference. At this stage, teams often rely on fuzzy matching to compare the incorrect entry against existing lot numbers.

Instead of looking for exact matches, fuzzy matching identifies close or probable matches for example, recognizing that “LOT-1298A” might actually be “LOT-1289A.” This helps narrow down the issue, but only after the problem has already disrupted operations.

Operations teams must then manually investigate receiving records, warehouse transfers, and production batches to confirm the correct lot. These investigations are time-consuming, often taking hours, and can delay production schedules

or shipments.

AI validation systems fundamentally change this process by applying fuzzy logic at the point of data entry. Rather than treating data as simply right or wrong, fuzzy logic evaluates whether an entry is likely correct based on context.

For instance, if a newly entered lot number doesn't exactly match supplier formats but closely resembles recent shipments, or if the quantity doesn't align with typical delivery patterns, the system assigns a confidence score instead of blindly accepting the input.

If the confidence is low, the system flags the entry immediately and prompts verification. This allows teams to correct errors in real time, before they propagate through inventory, production, and distribution systems.

In effect, fuzzy matching helps find errors after they occur, while fuzzy logic helps prevent them from entering the system in the first place.

Yield variance: The reality of physical ingredients

Food manufacturing introduces another complication that digital systems often struggle to handle: physical yield variance. Ingredients rarely reconcile perfectly with system quantities. Moisture loss, processing waste, and packaging differences can cause small variations between recorded and actual weights.

Traditional software often treats these differences as errors. AI systems can instead learn expected variance ranges and distinguish between normal operational variation and genuine data discrepancies. This reduces unnecessary alerts while still identifying real traceability problems.

THE AUDITOR WALKS IN.

You know the data exists.

But finding it? That's a different story.

For most food manufacturers, audit prep looks like this:

2–3 people pulled off the floor

Binders, spreadsheets & scattered email threads

45–90 minutes to respond to a single request

And a lot of hoping nothing is missing

That's not a data problem.

IT'S A SYSTEMS PROBLEM. LEARN HOW AI CAN HELP.

On May 6th at 11 AM, we're hosting a session for food manufacturing founders and CEOs — where you'll see exactly how to measure your audit exposure, and what to do about it.

No fluff. No theory.

Just the operational truth.



Reserve Your Seat
Know Where You Stand Before the Next Audit



03

The Co-Packer intelligence gap: How AI bridges the data divide without asking anyone to change

In the food and beverage industry, growth almost always involves Co-Packers. Emerging brands rarely own their first production facility. Instead, they partner with manufacturers who already have equipment, staff, and regulatory approvals.

This relationship is operationally efficient, but it creates a persistent challenge: data fragmentation. The Co-Packer's job is to manufacture product safely and efficiently. The brand's job is to track inventory, trace ingredients, and manage supply chain systems.

Those two responsibilities do not naturally produce the same data structures. For years, brands have tried to solve this gap by asking Co-Packers to

adopt new software platforms, change production documentation, or integrate with enterprise systems.

In reality, that strategy rarely works.

Most Co-Packers operate multiple production lines for multiple clients. Their processes are built for speed, safety, and compliance, not for adapting to every brand's technology stack.

AI-powered middleware is changing the equation. Instead of forcing Co-Packers to change how they work, these systems accept imperfect human-generated data & convert it into structured operational intelligence.

Why Co-Packers will always produce imperfect data

Even the most experienced Co-Packers operate in environments that make perfect data entry unrealistic.

Production floors are busy, noisy, and physically demanding. Operators often wear gloves, hairnets, and protective clothing, which makes interacting with digital interfaces difficult. Many production records are still written on batch sheets clipped to stainless steel workstations.

Replacing them entirely with digital entry can introduce new risks. In many facilities, a clipboard

and pen remain the most reliable method for documenting what happens on the floor.

This is why the Co-Packer intelligence gap exists.

Production data originates as human-generated records for accountability purposes, but brands require structured digital information for inventory, traceability, and financial systems.

The friction between these realities is where errors appear.



The accountability v/s automation balance

Food safety frameworks such as HACCP (Hazard Analysis and Critical Control Points) demand clear accountability. And there are products in the market that help with it.

Every production batch must have documented checks and operator sign-offs. These signatures

prove that safety procedures were followed.

Digital automation cannot fully replace that responsibility.

However, the data recorded during these checks must eventually enter systems used for:

Inventory management

Lot traceability

Regulatory reporting

Recall readiness


AI allows both worlds to coexist. Instead of forcing operators to enter perfect digital records during production, AI systems can interpret and validate information after it is captured, ensuring the data entering enterprise systems is accurate.

This preserves the accountability of handwritten records while preventing the data errors that traditionally occur when information is later transferred into software platforms.

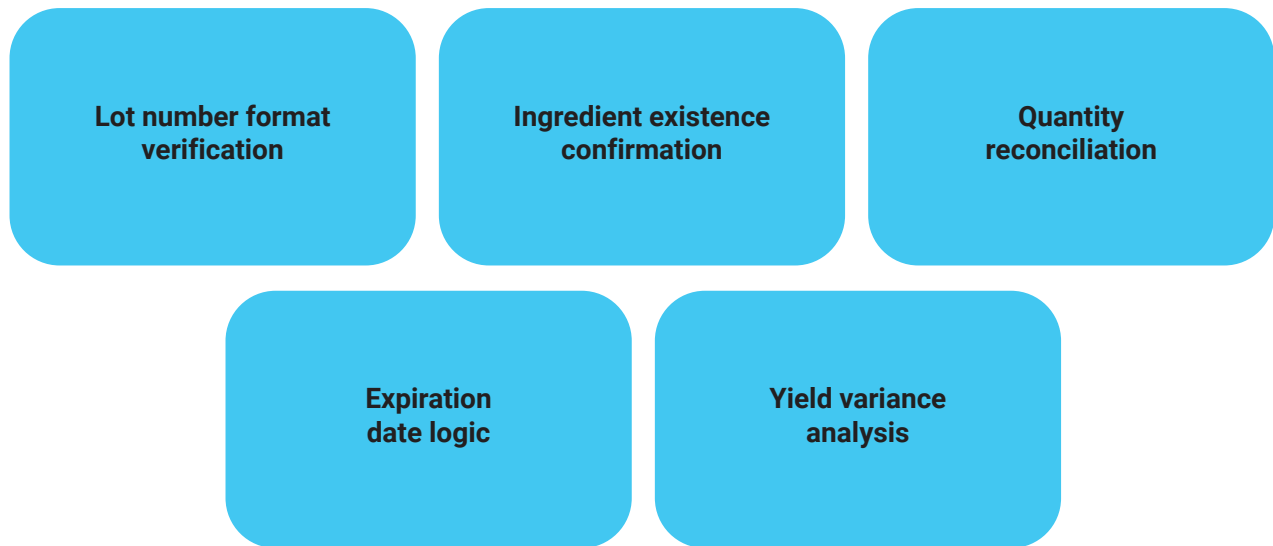
What an AI production submission layer looks like

In practice, an AI-powered middleware layer sits between the production floor and the brand's systems.

It is designed to be simple for the Co-Packer and intelligent for the brand. For the production worker, the interface might involve:



The system then performs automated validation checks before sending the data downstream. These checks include:



If something looks incorrect, the system flags it immediately rather than allowing the data to propagate through inventory and ERP platforms.

The key difference is that the system does not require the Co-Packer to change how they manufacture. It simply improves the reliability of the data generated during that process.

The most expensive Co-Packer data error

Among all production data problems, one stands out as the most common and costly:

Lot change communication failures. Production lines frequently switch ingredient lots during a run.

For example, a line producing protein bars may use:

Beef Lot A for the first half of the shift

Beef Lot B after the first pallet is depleted

If the batch sheet continues recording Lot A after the switch, the finished product records become inaccurate. This mistake is rarely intentional yet abundant. It happens because operators are focused on keeping the line moving. Front-office people don't know what's happening at the back. But the consequences can be serious.

During a recall investigation, incorrect lot assignments can force brands to expand the recall scope because the traceability chain is uncertain.

AI validation systems detect these errors by monitoring lot usage patterns and quantity flow.

If the recorded lot number cannot logically supply the quantity produced, the system flags the discrepancy immediately.

Instead of discovering the issue weeks later during an audit or investigation, the correction occurs while the production details are still fresh.



3.1

Validation rules that matter most

AI-powered production validation focuses on a small set of rules that capture the majority of real-world errors.




Lot existence verification:

The system checks whether a lot number recorded on the batch sheet actually exists in supplier or warehouse records. If a number does not match known inventory, it is flagged immediately.




Quantity reconciliation:

The system compares ingredient usage against recorded production volume. If the numbers fall outside expected ranges, the system alerts the operator or production supervisor.



Expiration date calculation:

Ingredients often have defined shelf lives. AI systems automatically verify whether the lot used during production was still valid at the time of manufacturing.



Yield variance tolerance:

No production process produces identical outputs every time. AI systems learn expected yield ranges and only trigger alerts when variance exceeds realistic limits. These checks sound simple, but together they capture the majority of operational data errors.

A real engagement example

In one engagement, **GrayCyan worked with a mid-sized beverage brand** to improve traceability and production data accuracy.

A recurring issue emerged: the same ingredient lot number appeared in batch records across multiple days, even though warehouse data showed the inventory had already been used. The cause was simple. An operator had copied the previous day's lot number while completing the batch sheet.

The error went unnoticed during production and was only caught later when the inventory system rejected the entries. By then, multiple records had to be reviewed, and operations were temporarily disrupted.

This is where **AI-driven validation changes the outcome**. By comparing ingredient usage against available lot-level inventory in real time, the system could have flagged the inconsistency immediately, preventing the error from spreading across multiple production cycles.



Key insight

The issue wasn't traceability, it was lack of real-time validation. Instead of relying on someone to catch the mistake later, AI ensures the system identifies what's wrong the moment it happens.

3.2

Meeting Co-Packers where they are

One of the most common mistakes brands make is assuming that their Co-Packers should operate like software platforms. They should not.

Co-packers exist to run production lines, manage quality controls, and ship finished product on schedule.

Expecting them to adapt their operations to every brand's data infrastructure is unrealistic. AI-powered middleware offers a different approach.

It meets Co-Packers where they already operate in physical production environments with human-generated documentation and converts

that information into structured, validated data that brands can rely on.

The result is not a perfectly digitized factory.

It is something far more useful:

Reliable operational data without forcing production teams to change how they work.

And for brands scaling through Co-Packers, that bridge between manufacturing reality and digital systems is becoming an essential part of modern food operations.



Game

Operational intelligence check

Can your system catch it before it costs you?

This isn't a quiz. It's a reflection of how your operations actually run.



?

The \$12,500 Deduction

A distributor submits a \$12,500 deduction tied to a promotion.
Your team assumes it's valid.

Three weeks later, someone realizes the promotion was never approved.

What happens in your system?

A.
Flagged instantly
based on promotion
mismatch

B.
Caught during
monthly
reconciliation

C.
Goes
unnoticed

Uncover the game's final secret—turn to the last page and claim the solution!

04

When your spreadsheet thinks: How AI is transforming the master order sheet into a living system

Every food and beverage brand eventually develops a spreadsheet that becomes far more important than anyone intended.

It rarely begins that way. At first, it's just a simple tracker. A founder creates a sheet to manage a few wholesale orders, track production runs, and ensure shipments leave on time. But as the business grows, the sheet grows with it.


New tabs appear, formulas multiply, status colors get added, and allocations start being tracked.

Eventually that file becomes something else entirely. It becomes the master order sheet.

Inside that spreadsheet lives the operational nervous system of the company: order status, distributor commitments, production schedules, inventory allocations, shipping timelines, and customer priorities.

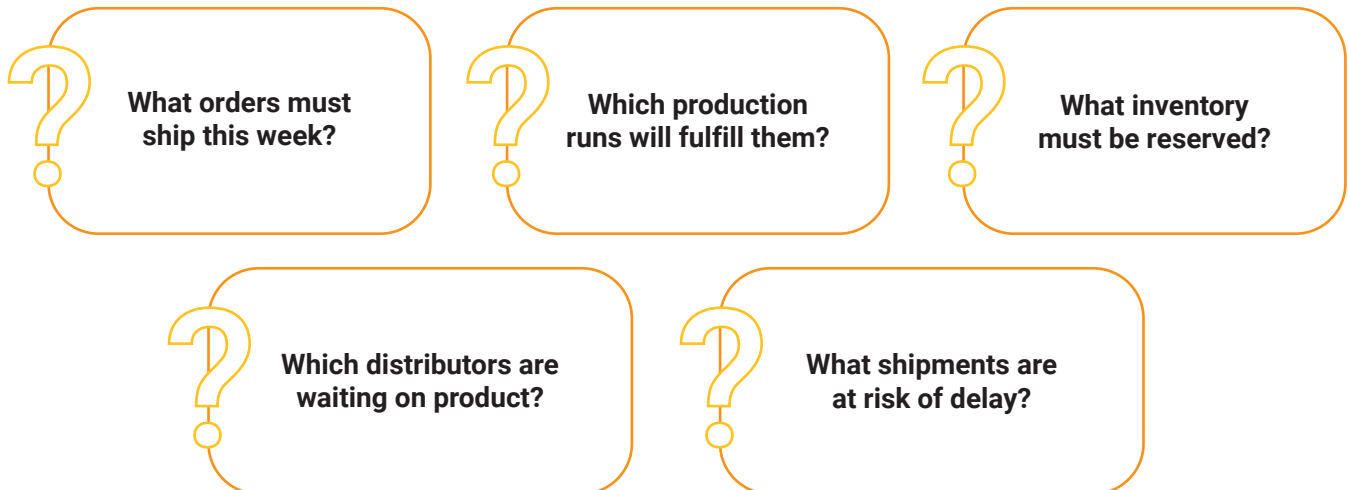
For many brands, this spreadsheet does things their ERP or inventory software never could. And yet, it is also one of the most fragile systems in the organization.

AI is not replacing the master order sheet. Instead, it is transforming it into something far more powerful: a living operational system that feeds itself real data.



The master order sheet exists because founders are solving real operational problems in real time. Most inventory platforms and ERP systems are designed to record transactions after they happen. But founders and operations leaders need tools that help them decide what should happen next.

A well-built master order sheet typically answers questions such as:



The sheet becomes a **visual command center**.

Instead of navigating multiple systems, teams see the operational picture in one place. That flexibility explains why founders trust it. The sheet reflects how the business actually operates, not how software designers assumed it would.

4.1 What spreadsheets do better than most systems

A thoughtfully designed master order sheet often outperforms enterprise systems in several key ways.

01

Visual operational clarity

Teams can immediately see which orders are waiting on production, which are ready to ship, and which are blocked by inventory shortages.

02

Cross-team communication

Sales, logistics, operations, and finance all rely on the same operational picture.

03

Real-time allocation decisions

When inventory is limited, teams can quickly decide which customers receive available product.

04

Operational flexibility

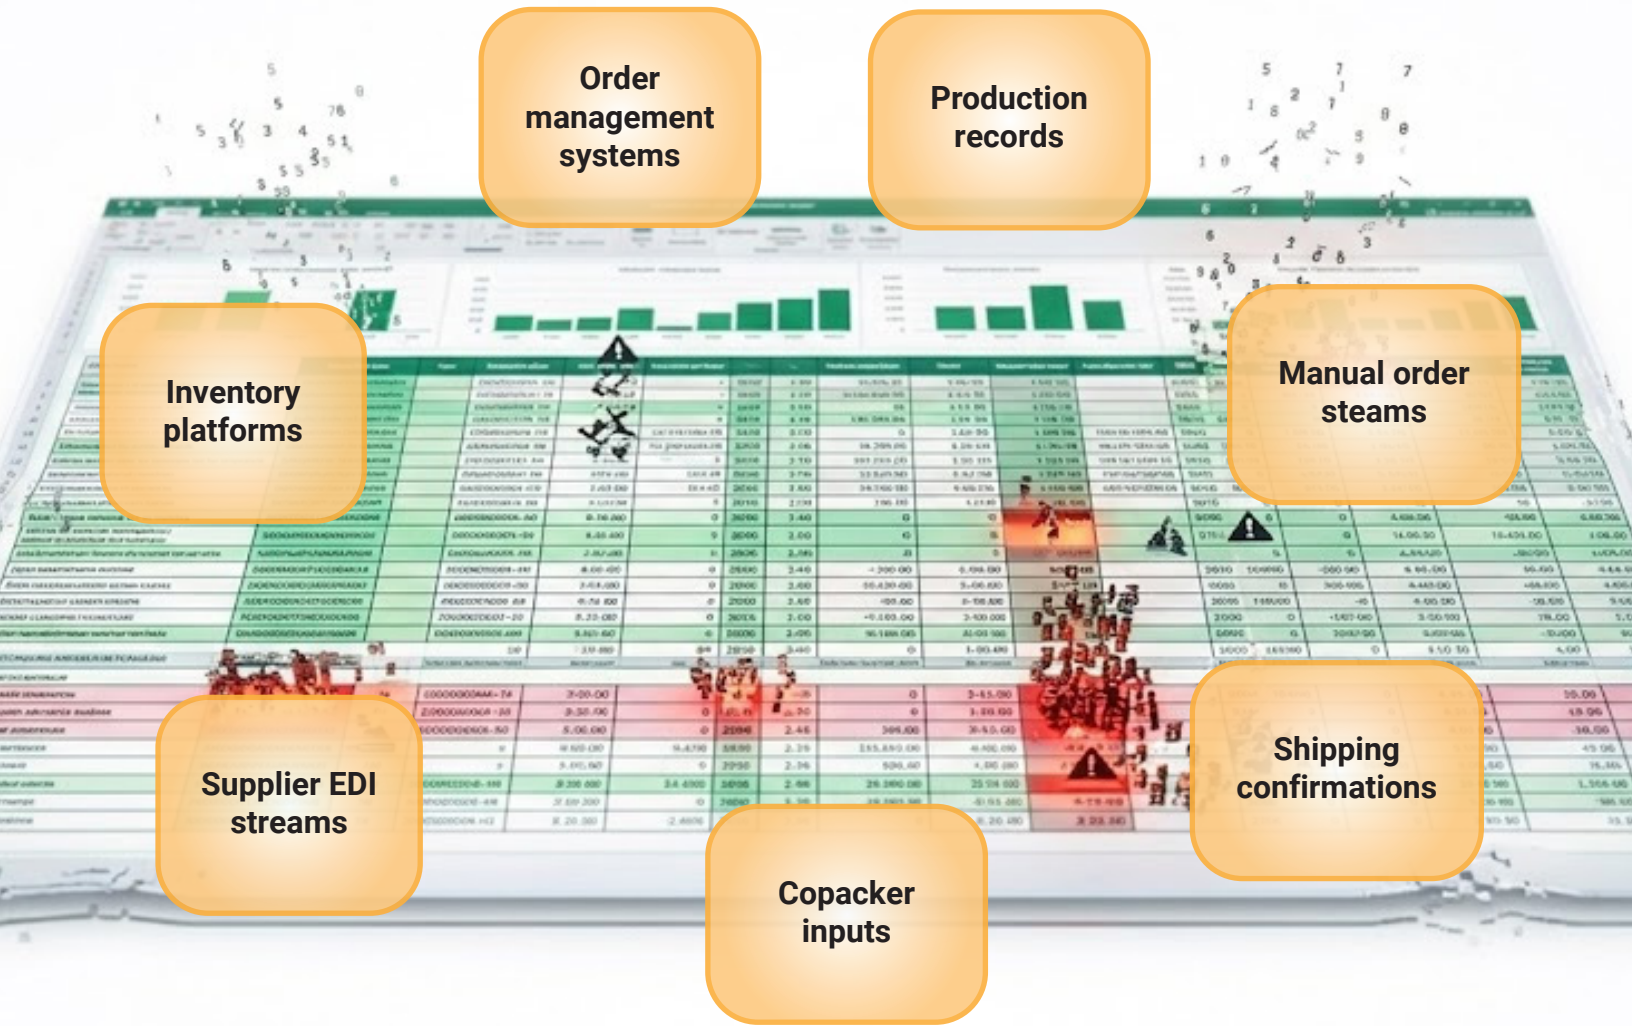
Columns can be added instantly. Rows can be reorganized. Teams adapt the sheet as the business evolves.

This is why many brands continue using spreadsheets even after implementing sophisticated software platforms.



Where the spreadsheet breaks

Despite its usefulness, the master order sheet has structural weaknesses. The most significant is manual maintenance. Every update requires someone to copy information from multiple systems:



Over time, the spreadsheet becomes dependent on one or two people who understand its logic. Without validation rules, errors can spread quietly.

There is usually no audit trail, no automated reconciliation, and no integration with the surrounding operational systems.

As the company grows, the sheet begins to behave like an operational system, but without the safeguards real systems require.

A GrayCyan example: When the sheet became the system



At GrayCyan, this pattern appears in almost every food and beverage client engagement. In one case, a fast-growing beverage brand was managing national distributor orders through a Google Sheet that had evolved over nearly three years. The sheet tracked:

Website and Amazon orders

Production scheduling

Shipping timelines

Distributor purchase orders

Warehouse inventory

Lot allocations

The founder described it as “the only place where the real status of the business lives.” But the sheet required constant manual updates.

Operations staff spent five to six hours every day copying information from their inventory system and logistics platform into the spreadsheet. When someone forgot to update a row, the entire operational picture became inaccurate.

GrayCyan implemented an AI-driven integration

layer that connected the sheet directly to the brand’s operational systems.

Production confirmations now update order status automatically. Inventory systems feed real-time stock levels into allocation calculations and shipment confirmations update delivery timelines, while the sheet itself remains unchanged and the familiar interface stays exactly the same. What had once been a fragile operational spreadsheet became a real-time command center.

SMB example: Sanzo and the reality of fragmented distribution

Sanzo, the Asian-inspired sparkling water brand founded by Sandro Roco in New York, provides a relevant example of how emerging food and beverage brands experience operational complexity as they scale into retail.

In a Shopify interview, Roco highlighted the fragmented nature of the grocery ecosystem, stating that it is: **“almost like a modern day miracle that our grocery stores offer the level of selection and curation [that they do] ... for how disparate the**

underlying systems are.”

Source: Shopify – Sanzo Food and Beverage Insights Interview

Sanzo has expanded its retail presence across major U.S. chains such as Whole Foods, Target, and Albertsons, reflecting the brand’s transition from early-stage operations into broader distribution.

Source: Retail Brew – Interview with Sandro Roco



Build a Successful CPG Brand with Sandro Roco

ing a ssful CPG Brand

Watch on YouTube

Scan to watch: How to build a successful CPG brand with Sandro Roco, founder & CEO at Sanzo

Roco’s statement points to a widely recognized characteristic of the grocery industry: **the presence of multiple, often disconnected systems across suppliers, distributors, and retailers.** This structural fragmentation is not unique to Sanzo but is a common feature of modern retail supply chains, where coordination occurs across numerous independent entities.

Industry interpretation: Operational reality for SMB brands

For many small and mid-sized (SMB) food and beverage brands operating within this environment:

Order data, inventory, and production planning are often managed across multiple platforms.

There is no single unified system connecting all stakeholders.

Teams must coordinate information across distributors, warehouses, and retail partners.

In practice, this often leads to the use of manual coordination tools, such as spreadsheets, to track and reconcile:



Purchase orders



Inventory levels



Production schedules



Fulfillment timelines

These tools act as a temporary operational control layer, helping teams bridge gaps between disconnected systems.

Operational constraint

However, such tools rely on:



Manual updates




Consistent human oversight



Version control across stakeholders

As businesses grow and orders increase, it becomes harder to manage everything manually, leading to delays, errors, and inconsistencies.



From fragmentation to systems thinking

This example aligns with a broader shift in the industry:

The challenge is not simply data availability, but the ability to integrate, standardize, and act on that data across systems. Emerging technologies, including AI-driven platforms, are increasingly being explored to address these gaps by enabling more structured and automated coordination across supply chains.

Sanzo's founder perspective highlights an important principle:

While consumers experience a seamless retail shelf, the systems behind it are often highly fragmented. For growing brands, the ability to move from manual coordination toward integrated systems becomes a critical factor in scaling operations effectively.

Connecting the spreadsheet to reality

AI changes the master order sheet in one crucial way.

It connects the spreadsheet to the operational systems around it.

Instead of relying on manual updates, the sheet becomes a live reflection of operational activity. Production confirmations update order status automatically. Inventory platforms feed current stock levels into allocation calculations.

Shipping systems update fulfillment timelines as soon as orders leave the warehouse. AI validation checks ensure the incoming data is accurate before it appears in the sheet. The result is something founders always hoped the spreadsheet would become: A real-time operational dashboard.

Once the master order sheet is connected to the systems around it, its role changes completely. What was once a manually maintained tracker becomes a live operational window into the business. Instead of relying on constant updates from the team, the sheet automatically pulls information from inventory platforms, production records, and shipment confirmations.

The structure remains familiar to the people who use it every day, but the data inside it is continuously refreshed, validated, and aligned with what is actually happening across operations. Rather than asking someone to manually refresh the sheet throughout the day, the system feeds it accurate operational data in real time.

The spreadsheet stops behaving like a fragile manual tool and starts functioning as an intelligent operational system that reflects the real state of the business.


A simple invitation

Many food and beverage brands already possess their most valuable operational tool. It's the spreadsheet they built themselves. The master order sheet often reveals more about how the business truly operates than any software system ever could.

The next step is simply connecting it to the systems around it.



Want to see how your operations can move from manual to intelligent?



4.2

When recalls become crisis: The role of data precision in containing risk

Not all recalls are equal. Two brands can face the same contamination issue on the same day. One resolves it quietly within a week. The other spends months dealing with retailers, regulators, and reputational fallout.




Source: FDA website

The difference between a contained recall and an expensive one often comes down to traceability precision. Stronger traceability supports faster identification and rapid removal of contaminated food. A 2025 study on overly broad recalls found median per-recall costs for producers ranging from \$3.0 million to \$72.7 million. Separate industry research has also put the average direct cost of a single recall at about \$10 million, excluding longer-term brand damage and lost sales.



Source: ScienceDirect website

That said, recall outcomes are shaped by multiple factors, including recall classification, regulatory requirements, legal guidance, and risk severity. Within those constraints, the level of traceability precision often determines how narrowly a recall can be defined and executed.



Why most recalls escalate

In theory, a recall is simple.

Identify the affected lot, trace where it went, and remove it from circulation.

Regulators increasingly expect companies to identify affected lots and produce traceability records very quickly. Under the FDA Food Traceability Rule, covered firms must be able to provide requested traceability information within 24 hours in an electronic sortable spreadsheet, and the rule is intended to enable faster identification and rapid removal of contaminated food from the market. In practice, most brands cannot answer three basic questions 'fast enough':

Source: FDA website



Which exact production run is affected?



Which units came from that run?



In many situations, where did those units end up?

In practice, the real test is speed: can the brand identify the affected lot, link it to the relevant production event, and show where it was shipped? FDA's traceability framework requires firms to maintain exactly those records and provide them within 24 hours.

When documentation is fragmented or inconsistent, brands are forced into approximation.

And approximation is expensive.

Instead of isolating a single lot, they recall multiple batches.

Instead of pulling from specific distributors, they notify many.

Instead of acting with confidence, they act with caution.



Source: FDA website

Instead of isolating the affected lot quickly, brands often have to broaden the recall scope. Instead of limiting action to clearly traced shipments, they may need to notify a wider set of consignees while records are being verified. That is exactly the problem FDA's traceability framework is meant to reduce by enabling faster identification and rapid removal through better lot-level records.

This is how a contained issue turns into an over-recall.

The anatomy of an over-recall

Over-recalls are rarely strategic decisions. They are defensive reactions to staggered data. Common failure points include:



Inconsistent lot naming conventions across internal and Co-Packer systems.



Manual data entry errors that break traceability chains.



Disconnected spreadsheets and systems that cannot reconcile production, inventory, and distribution fast enough (or even accurately in some cases).



Delayed visibility into where products actually moved after leaving the facility.

When these gaps exist, brands cannot confidently limit the scope. So they expand it. A 500-unit issue becomes a 5,000-unit recall not because risk increased, but because certainty decreased.

What precision looks like in practice

When documentation is structured and continuously validated, the recall process changes fundamentally. Instead of reconstructing events manually, the system can quickly identify:



The exact production window tied to the issue.



The specific Co-Packer batch and line involved.



The distribution path of each lot based on available shipment and partner data.



The last known movement of affected inventory across channels (noting exceptions where visibility is limited due to third-party constraints.)

This enables what can be described as a more targeted recall. Not everything produced that week. Not every distributor. Only what can be confidently traced based on available data. The operational impact is immediate, and the financial difference is substantial.

The cost gap no one plans for

The difference between a targeted recall and an over-recall is not incremental. It is exponential. A limited recall impacts logistics and inventory. A broad recall impacts the entire business. Costs typically include:



Large-scale product write-offs.



Reverse logistics and disposal expenses.



Retailer penalties and chargebacks.



Temporary delisting or loss of shelf space.



Long-term brand trust erosion.

In many cases, the largest cost is not the recall itself, but the loss of control over the narrative.

The Co-Packer blind spot

One of the most overlooked sources of traceability failure is the Co-Packer layer. Brands often maintain well-structured internal data, yet rely heavily on external manufacturing partners who:

Use different systems

Follow inconsistent data standards

Operate with varying levels of documentation rigor

This creates a critical gap. Even when internal systems are clean, the traceability chain can break at the point of integration. AI systems help address this, not by replacing Co-Packer workflows but by normalizing and reconciling external data against the brand's internal systems. This aligns with industry findings that interoperability and data standardization remain key barriers to effective traceability.



Sources: Deloitte & GS1 websites

The distributor blind spot (often the bigger risk)

An equally significant and often less discussed, blind spot lies in the distribution layer.

Modern supply chains rely on multiple intermediaries, including distributors, logistics providers, and retail networks. However, end-to-end visibility across these networks is often limited, particularly once products move beyond the manufacturer's direct control.

In practice, brands may know what was shipped, but not precisely:

Where each unit currently resides

How inventory is redistributed across locations

Which specific regions or channels received affected batches

This lack of downstream visibility creates a structural limitation. Once products move across multiple supply chain partners, traceability becomes fragmented and increasingly difficult to maintain at a granular level.

Industry research supports this, highlighting that limited visibility across supply chain partners remains a major challenge for achieving true end-to-end traceability.



Source: GS1 website



Where AI still creates leverage

Even when perfect traceability is not achievable, AI can still provide meaningful operational advantage. Rather than relying solely on deterministic tracking, AI enables probability-driven decision-making, including:



Pattern-based inference:

Using historical shipment and sales data to estimate likely product locations



Risk clustering:

Identifying high-probability regions or distribution nodes



Recall prioritization:

Focusing on the most likely impacted zones first



Scenario modeling:

Simulating product flow across networks

This reflects a broader shift toward advanced analytics and AI to improve supply chain decision-making under uncertainty.

Source: Mckinsey website





The regulatory pressure is rising

Regulators are increasingly pushing for faster and more precise traceability.

Under the FDA’s Food Traceability Rule (FSMA 204), companies must provide required traceability records within 24 hours during a recall or investigation.



Sources: Deloitte & GS1 websites

However, while these regulations improve traceability requirements, they do not fully eliminate visibility gaps across complex, multi-party distribution systems, as the rule focuses on recordkeeping rather than real-time tracking.

The reality: Visibility has limits

AI significantly enhances traceability, but it does not eliminate structural blind spots.


Co-packer gaps require better integration and standardized data exchange

Distribution gaps require managing complexity across multiple independent partners

This aligns with broader industry consensus that data fragmentation across supply chain ecosystems remains a core challenge. The future of traceability won’t be defined by perfect visibility—it will be defined by how intelligently companies operate in the presence of complexity and incomplete data.

What strong recall execution signals

Retailers and distributors evaluate recalls differently than brands expect. The issue itself matters. But the response matters more. When a brand demonstrates:



Fast identification of affected units



Clear and limited recall scope



Accurate communication backed by data

It signals operational maturity. In contrast, delayed responses and broad recalls signal risk. Over time, this distinction affects:

Shelf placement decision

Distribution partnerships

Willingness to onboard new SKUs

Moreover, recall execution becomes a proxy for operational reliability. AI does not eliminate recalls, but it transforms them from reactive, uncertain, and disruptive events into controlled, precise, and contained operations. This change reshapes how businesses manage risk and respond under pressure.



05

AI on the floor: What's working, what's not, and what every food & beverage operator needs to know right now

Artificial intelligence is being marketed to food and beverage companies as a solution that can sit on top of broken systems and magically make operations smarter.

It cannot.

AI does not fix bad processes. However, it doesn't mean operators go about fixing processes in an 'ideal way' because it doesn't exist.

The reality inside many food and beverage operations is simple: the systems are fragmented, the data is messy, and the operational picture is assembled manually.

Before AI can optimize forecasting, production planning, or supply chain decisions, something far

more basic must exist: clean, reliable data at the source.

Most operators assume they have it.

The companies seeing real value from AI today are not the ones deploying flashy predictive tools. They are the ones using AI to solve small but expensive operational problems that occur every day on the production floor.

More importantly, they are building systems around how their processes actually work, rather than forcing processes to fit disconnected tools. When technology supports real workflows, data becomes cleaner, decisions become faster, and operations run more efficiently.

5.1 Where AI is actually delivering ROI today

Despite the hype surrounding artificial intelligence, only a handful of applications are consistently delivering measurable operational value across food and beverage companies today.



1. Error prevention at data entry

The most practical use of AI in food operations today is preventing incorrect data from entering the system in the first place. Many inventory and production systems accept whatever operators enter.

Lot numbers can be mistyped, quantities can be entered incorrectly, and expiration dates can be recorded inaccurately. These errors may not appear until days later when systems reject transactions or inventory discrepancies emerge. AI validation layers change that dynamic.

Instead of blindly accepting inputs, the system checks them immediately against known rules:

Does the lot number exist?

Does the quantity reconcile with production volume?

Does the expiration date make sense based on the ingredient's shelf life?

If something looks wrong, the system flags it instantly. Correcting the error takes seconds instead of hours of investigation later. For many companies, eliminating these small operational mistakes produces immediate time savings and more reliable inventory data.

2. Automated transaction creation in inventory systems from production data

Another area where AI is delivering strong operational value is automating inventory transactions based on confirmed production activity. In many food companies, production teams record completed batches on paper or simple digital logs. Operations teams later translate those records into system transactions:



This translation process can take hours each week. AI systems can read confirmed production records and automatically generate the correct transactions in inventory systems. Once a production batch is verified, the inventory system records ingredient usage and finished goods creation without manual entry.

This removes repetitive administrative work and reduces the risk of transcription errors.

3. Inventory anomaly detection

The third area where AI consistently delivers value is detecting inventory anomalies before they become operational problems.

Inventory discrepancies are common in food operations. They occur when ingredient consumption does not match recorded production activity or when shipment records do not align with warehouse inventory. Traditionally, these issues are discovered only when someone notices a problem, often days or weeks later. AI systems continuously monitor inventory patterns and flag inconsistencies early. For example:



Instead of relying on periodic inventory checks, operators receive alerts as soon as anomalies appear.



Why AI forecasting often fails

Many food and beverage companies begin their AI journey with forecasting or demand planning tools. These platforms promise better inventory planning, improved production scheduling, and more accurate demand predictions.

But predictive models rely on historical data patterns. If the underlying operational data is inconsistent, the model cannot produce reliable insights. Common problems include:



Inventory corrections that distort historical stock levels



Missing production records



Incorrect shipment confirmations



Unrecorded ingredient substitutions

When this type of data feeds an AI forecasting model, the predictions may appear mathematically sophisticated, but they are built on unreliable information.

In simple terms: **Garbage In Garbage Out.**

Before investing in forecasting AI, operators must first ensure that operational records accurately reflect reality.

5.2 The role of AI middleware

The most practical AI architecture in food and beverage operations today is not a standalone platform. It is middleware. Middleware sits between existing systems, such as:




Its job is to validate, reconcile, and translate the data moving between them. Many tasks currently performed manually by operations teams happen inside this layer:



Instead of forcing companies to replace their existing systems, middleware improves how those systems communicate.

This is why it has become one of the fastest-growing approaches to applying AI in operational environments.



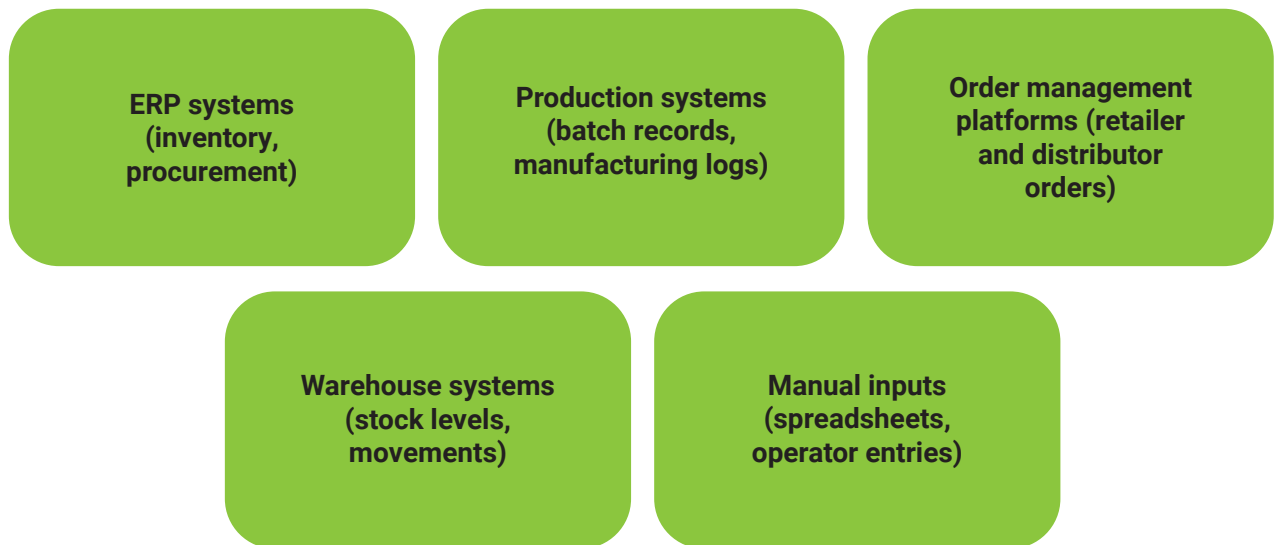
Five questions to ask any AI vendor

Food and beverage operators evaluating AI tools should approach vendor claims with healthy skepticism. Before signing a contract, ask five practical questions.


Where does the system get its data?

Every AI system is only as reliable as the data it consumes. Before implementation, it's critical to map out all data sources feeding the system.

In food & beverage and processing operations, this typically includes:



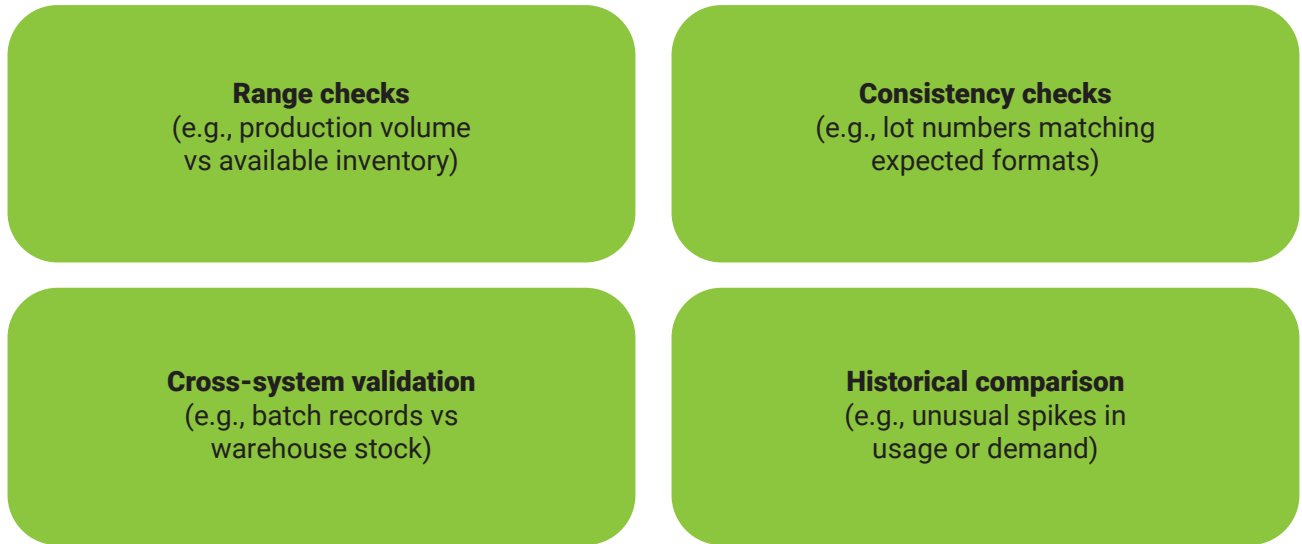
The key questions are:




If data is fragmented, delayed, or duplicated, the AI system will reflect those inconsistencies, just faster. Clean inputs are not optional; they are foundational.

How does the system validate incoming data?


AI systems should not passively accept data, they should actively question it. Effective validation includes:



Without validation, incorrect data flows downstream and compounds across:



The result is not just bad data, it's bad decisions at scale.



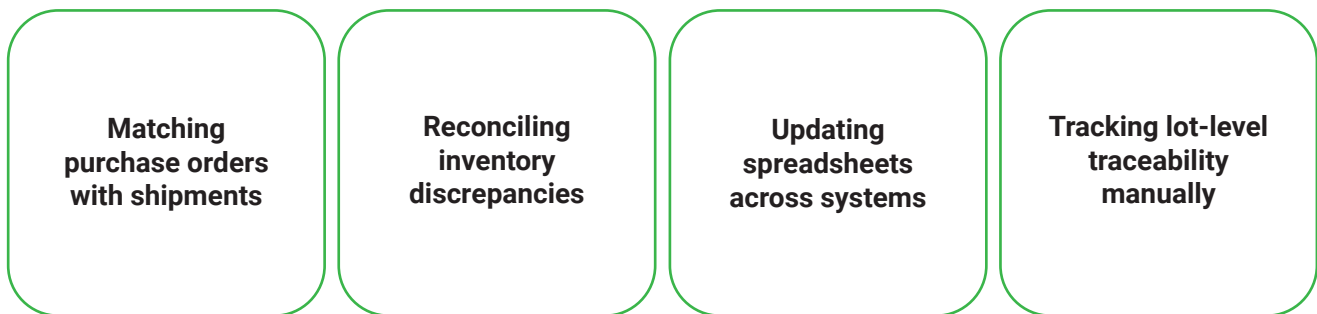


What manual work will this system eliminate?

AI should solve real operational friction, not just add another layer of technology. Identify:



Examples include:



If the system does not clearly reduce manual workload, it may improve reporting—but not operations.

What happens when the system encounters incorrect data?

No system operates in a perfect environment. The real test is how it handles errors. A well-designed AI system should:

Flag inconsistencies immediately

Provide clear explanations of what's wrong

Allow users to correct data at the source

Prevent invalid data from moving downstream


A weak system will:

Accept incorrect inputs silently

Surface issues too late

Require manual investigation after the fact

The goal is not just automation it is controlled, explainable automation.



How quickly will the operational team see measurable improvements?

Operational teams adopt systems that deliver results, not promises. A practical AI system should:



Look for measurable indicators such as:



If value is delayed or unclear, adoption will slow, and the system risks becoming another unused tool. Choosing the right AI partner is less about the sophistication of the technology and more about the practical value it delivers inside daily operations. Food and beverage companies operate in environments where small data errors can cascade into production delays, inventory discrepancies, and costly investigations.

The right AI solution should simplify those realities, not add another layer of complexity. By asking the right questions early, founders can separate genuine operational tools from marketing promises and ensure that any AI investment strengthens the systems already running their business rather than disrupting them.

Artificial intelligence is not a magic layer that fixes broken operations. But it can solve the small operational problems that quietly consume time, create errors, and obscure the true state of the business.

Game

AI reality check

Is your operation ready for AI or just buying into it?

Every food and beverage company says they are “using AI.”
Very few are actually getting value from it.

Answer honestly.



The Forecasting Trap

An AI vendor promises 95% demand forecasting accuracy.
Your data includes:

**Manual
inventory
corrections**

**Missing
production
records**

**Inconsistent
shipment
confirmations**

What should you do?

A.
Implement the
forecasting tool
immediately

B.
Clean and validate
operational data
first

C.
Use the tool
and manually
adjust outputs

Uncover the game's final secret—turn to the last page and claim the solution!



06

Winning back the deduction: How AI-powered documentation is turning distributor disputes in favor of small brands

For many growing food and beverage brands, distributor deductions are an unavoidable part of doing business. What makes them particularly frustrating is that many of them occur even when the brand shipped the order correctly.

A distributor may claim the shipment was short, report missing cases, or flag incorrect labeling or documentation and the result is simple: the distributor short-pays the invoice.

For large brands, these deductions are annoying but

manageable. They have teams, systems, and documentation infrastructure designed to dispute them.

For small and mid-size brands, deductions often go uncontested. Not because the distributor is right but because proving otherwise takes time and documentation most brands do not have readily available.


AI-powered documentation systems are beginning to change that dynamic.

6.1

How distributor deductions actually work

In a typical distribution relationship, a brand ships product to a distributor and sends an invoice; if the distributor believes there is an issue such as missing cases, damaged goods, or incorrect paperwork they issue a deduction and pay less than the invoiced amount.

Distributor deductions can happen for several reasons, such as short shipment claims, transit damage, incorrect labeling or packaging, missing paperwork, or late deliveries. Some are legitimate, but others result from warehouse errors or miscommunication. In most cases, the burden of proof falls on the brand, and if the brand cannot quickly produce documentation showing what was shipped, the deduction usually stands.



Why small brands lose most deduction disputes

Large CPG companies typically use sophisticated systems that automatically capture and organize shipment documentation. Smaller brands, however, rarely have that infrastructure in place.

Instead, the information needed to dispute a deduction is often scattered across multiple locations shipping confirmations in one system, packing lists in another, bills of lading stored as PDFs, warehouse photos saved on a phone, and lot-level shipment records recorded manually.

In 2026, Glimpse announced a partnership with PLTFRM to bring its AI-driven deductions platform to CPG brands. The announcement followed PLTFRM's 2025 acquisition of Presence Marketing, which expanded PLTFRM's reach across the CPG sector. Glimpse says its platform helps brands

recover lost revenue, automate deductions workflows, and reconcile chargebacks faster. In separate company materials, Glimpse states that many CPG brands lose 2–3% of annual revenue to invalid deductions. Together, those claims indicate that deduction recovery has become a large enough operational problem to support specialist AI platforms built specifically for it.

When a deduction appears weeks later, the operations team must reconstruct the entire shipment history. Tracking down the necessary records can take hours, sometimes longer. Because small teams are already focused on production, fulfillment, and sales, many of these disputes are never pursued, and the brand simply absorbs the loss.

The documentation chain that wins disputes

Successful deduction disputes depend on a clear and complete chain of evidence. The key documents typically include packing lists showing what was shipped, lot-level shipment records identifying which product lots were included, signed bills of lading confirming carrier pickup, warehouse receiving confirmations verifying delivery, and photographic evidence documenting the condition of pallets at the time of shipment.

Together, these records establish a reliable timeline of what left the warehouse and in what condition. The real challenge is not generating these documents most brands already have them but locating and assembling them quickly when a dispute arises.

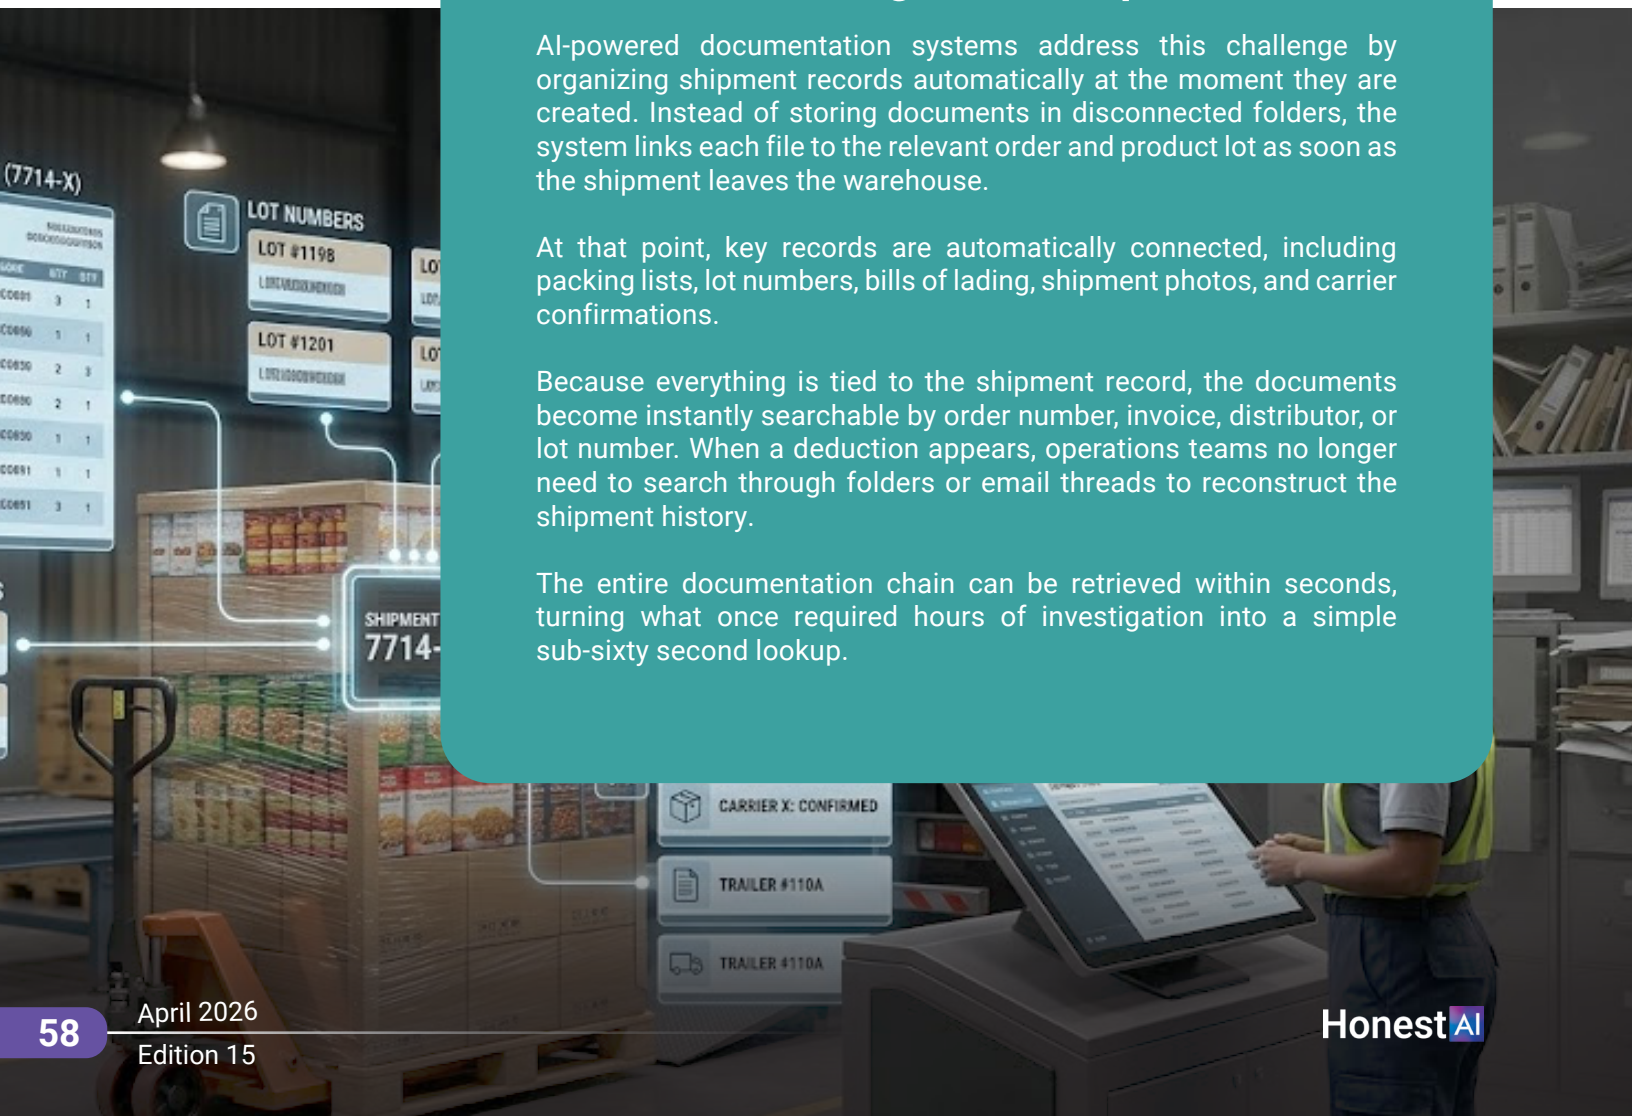
Where AI changes the equation

AI-powered documentation systems address this challenge by organizing shipment records automatically at the moment they are created. Instead of storing documents in disconnected folders, the system links each file to the relevant order and product lot as soon as the shipment leaves the warehouse.

At that point, key records are automatically connected, including packing lists, lot numbers, bills of lading, shipment photos, and carrier confirmations.

Because everything is tied to the shipment record, the documents become instantly searchable by order number, invoice, distributor, or lot number. When a deduction appears, operations teams no longer need to search through folders or email threads to reconstruct the shipment history.

The entire documentation chain can be retrieved within seconds, turning what once required hours of investigation into a simple sub-sixty second lookup.



6.2

The financial impact of deduction recovery


Distributor deductions can quietly erode margins, particularly for brands scaling through retail distribution.

The financial impact grows quickly as revenue increases. A brand generating around \$5 million in annual distribution revenue can already have tens of thousands of dollars tied up in deduction disputes.

At \$10 million, the financial risk becomes much more significant.

By \$20 million, unresolved deductions can add up to hundreds of thousands of dollars in potential losses.

Recovering even a portion of those claims can materially improve a brand's profitability. The brands that consistently win deduction disputes are rarely the ones with perfect shipments.




Advice from the operations side

Operators who have dealt with deduction disputes know a simple truth: these arguments are rarely won through explanation, they are won through documentation. The faster a brand can produce a complete shipment record, the stronger its position becomes.

The solution isn't to reinvent operations or overhaul existing processes. It's to have a system that knows exactly where each piece of information resides and can bring it together instantly.

AI does not eliminate deductions entirely, but it shifts the balance of power. By organizing data across systems and making it searchable and accessible in real time, AI enables brands to respond quickly and confidently.

For growing companies, that shift transforms deductions from a frustrating financial drain into a manageable and defensible operational process.



Industry recognized AI maturity model

As deduction volumes grow and documentation becomes harder to manage, the real question isn't whether you need AI, it's whether your current systems are ready for it. An AI maturity assessment helps identify how well your data, workflows, and documentation processes are structured to support automation, real-time retrieval, and audit readiness. It highlights gaps in visibility, standardization, and system integration, areas that directly impact your ability to recover deductions efficiently.

For brands scaling through retail and distribution, understanding this maturity level is the first step toward building systems that don't just store data, but actively support faster decisions, stronger disputes, and more predictable outcomes.

Assess your operational AI maturity today with GrayCyan and take control of your deduction recovery process.



AUTOMATE SOP CREATION FROM VIDEOS

Teams waste hours manually creating process manuals. Kaizenify eliminates that friction.

**Streamline operations.
Standardize faster.
Scale smarter.**



Book a Demo



Upload a process video and instantly generate:

1

SOP titles

2

Step-by-step workflows

3

Visual documentation

4

Time estimates



07

Scaling the kitchen: How AI helps scale beyond your first Co-Packer

For most food and beverage brands, the first Co-Packer relationship feels like a milestone. Production moves out of the founder's kitchen and into a professional facility. Capacity increases. Quality improves. The business finally begins to scale.

But the second Co-Packer is where real operational complexity begins.

At a certain point, usually when demand grows faster than a single facility can support brands must add another manufacturing partner. It might be to increase volume, reduce shipping costs by producing closer to regional markets, or mitigate supply chain risk. In theory, it's a natural step in

growth.

In practice, it often becomes a logistical nightmare. The reason is simple: most emerging brands unknowingly build their entire operation around a single Co-Packer. Their processes, spreadsheets, inventory tracking, lot traceability, and production schedules all evolve around that one partner's systems and workflows. When a second facility enters the picture, the system that once worked suddenly breaks.

Instead of doubling capacity, the business doubles complexity. This is where AI is quietly transforming how scaling works.



The hidden operational requirements of a second Co-Packer

Most founders assume that adding a second production partner is primarily a legal and contractual exercise negotiating pricing, confirming certifications, and aligning production timelines.

But the real challenge appears after the contract is signed. Suddenly, the business must reconcile:

Two different production reporting formats.

Two separate communication and escalation workflows.

Two different supplier coordination and procurement alignments.

Two scheduling structures.

Two quality assurance documentation processes.

Even something as simple as a production report may arrive in completely different formats - one Co-Packer sending structured ERP exports, while another delivers manually prepared spreadsheets.

Without a unified system, teams spend hours reconciling data manually. Errors creep in. Traceability becomes harder. Operational visibility disappears. This is why many brands discover too late, that their infrastructure isn't built for multi-partner production.

The role of AI as production middleware

The breakthrough enabling multi-Co-Packer scaling isn't robotics or predictive analytics. It's something less visible but far more impactful: AI-driven data normalization.

Think of it as middleware between the brand and its production partners.

Instead of forcing Co-Packers to adopt new software, something that rarely succeeds AI systems ingest production data in whatever format the facility provides. Machine learning models then classify, interpret, and translate that data into standardized internal records.

In other words, the AI acts as a universal translator for manufacturing data. A production log from Co-Packer A and a spreadsheet from Co-Packer B can both be converted into the same structured dataset:



Standardized lot numbers



Unified ingredient traceability



Consistent inventory updates



Normalized production reporting

The result is a production-agnostic data layer, a system where the brand's operational intelligence exists independently of any individual Co-Packer's technology stack.

This architectural shift is subtle but powerful.

Instead of asking Co-Packers to adapt to the brand's systems, the brand builds systems flexible enough to accommodate any Co-Packer.

Why certified Co-Packers still create data chaos

Ironically, many of the best manufacturing partners, those with strong food safety credentials introduce the greatest operational complexity.

Facilities certified under standards like SQF, BRCGS, or FSSC 22000 often run highly specialized internal systems designed for regulatory compliance. These systems generate detailed production documentation, but they rarely align with the reporting structures used by emerging brands.

For example, a Co-Packer may track lot genealogy across multiple ingredient batches and production

runs using internal codes that mean little outside their facility. When those reports reach the brand's operations team, translating them into actionable inventory data can become a manual process.

AI systems solve this by learning the mapping between the facility's internal identifiers and the brand's operational data model.

Over time, the system becomes increasingly accurate at interpreting each Co-Packer's reporting style, automatically reconciling data that once required human intervention.

The timeline difference: AI v/s traditional scaling

The operational timeline for adding a second Co-Packer differs dramatically depending on whether AI infrastructure is already in place.

Without AI support, onboarding a new production partner can take months of manual system alignment. Teams must redesign spreadsheets, build new reporting templates, train staff on reconciliation processes, and troubleshoot inconsistencies across datasets.

With AI middleware, much of that complexity reduces.

Instead of redesigning internal workflows, the AI layer simply learns how to interpret the new Co-Packer's data formats. Production visibility remains consistent from day one.

What previously required months of operational adjustments can happen in weeks.

Preparing before the next Co-Packer

The most forward-thinking food brands aren't waiting until the second Co-Packer is signed. They're building the infrastructure before the need becomes urgent.

Consider a fast-growing functional beverage brand preparing to expand production into a second region next year. Rather than waiting for that moment, the company has already begun implementing an AI-driven operational layer that standardizes production data across partners.

By the time negotiations with the new Co-Packer are finalized, the brand's internal systems will already be designed to accept the incoming production data. In effect, the company is building the operational scaffolding for growth before the additional capacity arrives.



The scaling mantra for food manufacturers

Historically, scaling food production meant securing more manufacturing capacity. Today, scaling means something different. It means building operational intelligence that can function across multiple production partners simultaneously. AI is making that possible by removing the data friction that once limited growth.

The brands that understand this shift early won't just add Co-Packers more smoothly, they'll scale with far greater resilience.

And in a food industry increasingly defined by distributed manufacturing, that capability is quickly becoming a competitive advantage.

08

The blueprint before the build: Why the most important AI investment you'll make is documentation

Artificial intelligence has become the most discussed technology in modern business. From predictive analytics to automated inventory forecasting, the promise of AI in food and beverage operations is immense.

Yet an uncomfortable truth sits beneath many AI initiatives. Most fail before they begin. Not because the algorithms are flawed. Not because the software is inadequate.

But because the operation itself was never properly documented.

Across food and beverage businesses from

fast-growing startups to established brands, the same pattern appears repeatedly. Teams attempt to automate workflows that have never been clearly defined. Integrations are assumed rather than specified. Business rules exist only in the memories of a few key employees.

When consultants begin building AI systems on top of this ambiguity, the project inevitably stalls. Edge cases appear. Exceptions multiply. The cost of correcting assumptions mid-build becomes enormous.

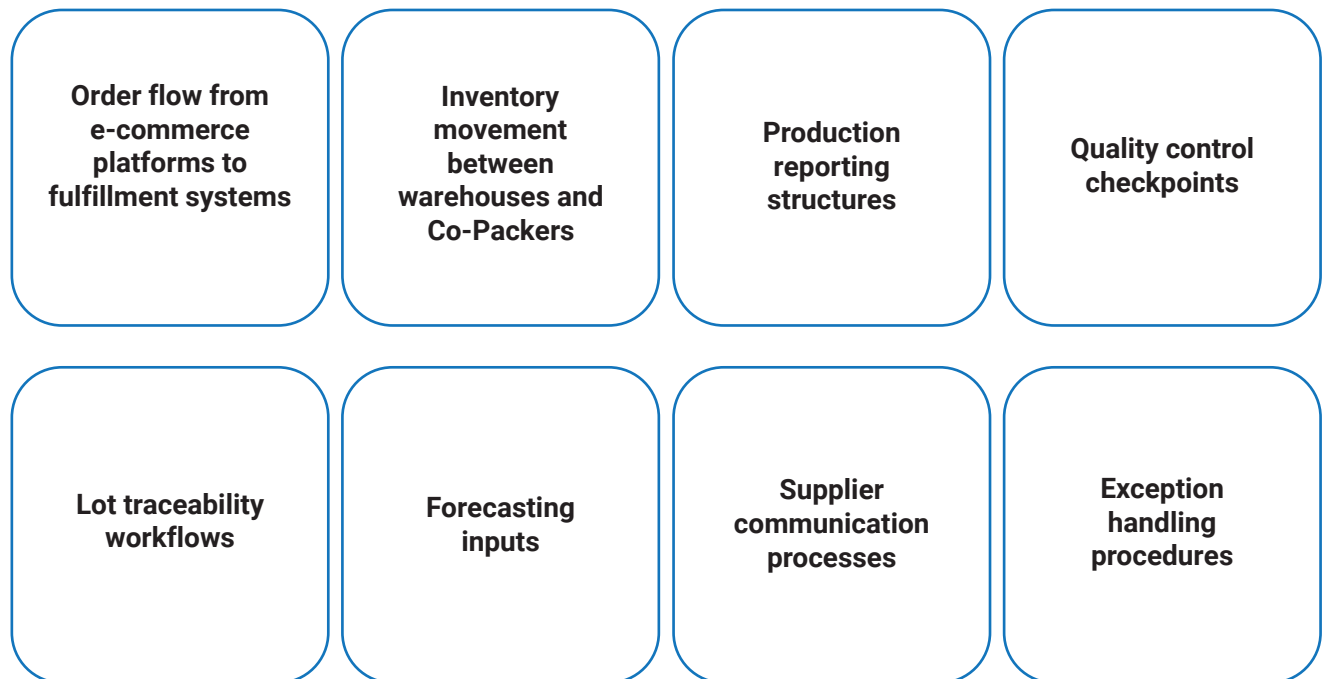
The problem isn't the AI. The problem is the missing blueprint.

The operational manual most businesses never wrote

Most businesses focus heavily on sales, marketing, and growth strategies, but very few take the time to document how their business actually runs on a day-to-day basis. Processes exist in people's heads, decisions happen informally, and critical workflows are often managed through habit rather than structure. While this may work in the early stages, the absence of a clear operational manual eventually leads to inefficiencies, inconsistent execution, and heavy dependence on specific individuals within the organization.

An operational SOP serves as the backbone of a scalable business. It brings clarity to processes, defines responsibilities, and ensures that tasks are executed consistently regardless of who is performing them. When properly documented, it allows teams to work more independently, improves accountability, and makes onboarding and growth significantly smoother. Despite its strategic importance, it remains absent in many organizations, making it one of the most vital documents that rarely exists.

In a food and beverage business, this blueprint might include:



It also includes the workarounds, the small manual fixes employees perform daily to keep systems aligned. These details are often invisible until someone attempts to automate them.

Why the blueprint is worth more than the AI

One of the most surprising discoveries companies make during this process is that the blueprint itself becomes a strategic asset. It becomes the operational manual the business never had.


New hires can understand workflows immediately

Consultants can build systems without guesswork

Leadership can identify inefficiencies that were previously hidden

Perhaps most importantly, the business becomes less dependent on any individual employee's institutional knowledge.

For companies preparing for acquisition or investment, this level of operational clarity is extremely valuable. Buyers and investors consistently evaluate whether a business can operate independently of its founders or early team members.




Phase 1: The step most AI projects skip

In successful AI engagements, the first phase rarely involves building software. Instead, it involves discovery and documentation.

During this stage, teams map the entire operational architecture of the business. Systems are catalogued. Data flows are identified. Business rules are extracted from the people who currently manage them.

Every assumption is tested. This phase may feel slow compared to the excitement of launching AI tools, but it dramatically reduces risk later in the project. Brands that skip this step often discover they must repeat it mid-build, after system builds have already begun. At that point, revisions become expensive. In other words, skipping the blueprint doesn't save time. It simply postpones the work to the most costly moment.



Industry recognized AI maturity test

There is a simple test that reveals whether a business is ready to implement AI.

Ask one question:

If the person who understands your operation best left tomorrow, could an AI system be trained using only what is written down?

If the answer is no, the business isn't ready for AI.

This test exposes the gap between operational knowledge and operational documentation. Many companies assume their processes are clear, until they attempt to describe them in precise, repeatable terms. AI systems require that precision.





Build **AI-Ready** Operations with GrayCyan

GrayCyan works with CPG brands and food processors to transform complex, fragmented operations into structured, AI-ready systems. These systems are built from the ground up—based on real workflows, not assumptions.

What sets GrayCyan apart:

On-site factory immersion

We observe real workflows, not just documented ones.

Job shadowing your team

Capturing how work actually gets done.

Workflow standardization

Turning tribal knowledge into scalable systems.

AI-readiness blueprints


Designed for automation, visibility & growth.

Audit readiness

Building AI first audit readiness.

The outcome: Clear, structured operations that enable automation, reduce inefficiencies, and prepare your business for AI-driven scale.

Ready to turn your operations into a competitive advantage? Let GrayCyan map your path to AI readiness.



What the first 90 Days of an AI engagement look like

For small and mid-sized food businesses, a serious AI initiative typically unfolds in three phases during the first ninety days.

The first phase is discovery. Teams interview stakeholders, review systems, and observe workflows to understand how the business truly operates, not how it is assumed to operate.

The second phase is blueprint development. All operational knowledge is translated into structured documentation that maps systems, integrations, and decision rules.

The third phase is targeted build. Only after the blueprint exists do developers begin constructing the first AI-supported workflow usually the area with the highest immediate leverage, such as inventory forecasting, production reconciliation, or demand planning.

This sequence dramatically increases the likelihood of success.

Instead of building AI in the dark, the system is built on a clear operational foundation.

The first step toward AI transformation

The most important insight for food and beverage founders may be this:

AI transformation doesn't start with technology.

The brands that succeed with AI are rarely the ones with the largest budgets. They are the ones that understand their operations deeply enough to document them.

Once that blueprint exists, automation becomes possible.

Without it, even the most advanced AI platform will struggle.

The next step

For companies curious about where AI could create the most operational leverage, the first step isn't a software demo.

GrayCyan offers a focused, 1-month AI consultation designed specifically for food and beverage businesses ready to move beyond theory and achieve real, measurable impact.

During this engagement, GrayCyan takes a deep, structured look at your operations, identifying exactly where AI can drive immediate, measurable value. From there, Phase 1 of the AI implementation is clearly defined, focusing on practical outcomes.

This isn't about experimenting with tools or following trends. It's about creating a clear, execution-ready blueprint because before investing in new systems, you need a strategy that ensures they work

If you're looking to streamline operations, improve traceability, or strengthen compliance readiness, this is where your AI journey starts, with clarity and direction.

As a next step, join GrayCyan's upcoming webinar to see how AI can seamlessly connect your systems and elevate USDA audit readiness into a strategic strength.



Conclusion:

As AI continues to transform industries, the conversations around it must evolve just as rapidly.

Honest AI exists to cut through the noise not to sell, not to persuade, but to question, clarify, and bring forward what truly matters.

In a landscape filled with bold claims and opaque systems, transparency is no longer a luxury, it is the foundation of trust. And trust will ultimately define whether AI becomes a force for empowerment or a source of uncertainty.

This magazine is a step toward that clarity, a space where ideas are explored with intent, assumptions are challenged with rigor, and real-world impact takes priority over abstraction.

Building better AI is not just about improving models. It is about understanding how they are used and where they truly matter.

The future of AI will be shaped by those who question it, challenge it, and take responsibility for it.

We hope this edition gave you a fresh perspective.

The next one is coming soon.

Validate your answers

Game 1:


Answer = A Flagged instantly based on promotion mismatch

Reason = If your system flags the deduction instantly based on promotion mismatch, it means your operations are connected and validated in real time. B and C indicate delayed or missing control, which leads to revenue leakage.

Game 2:

Answer = B Clean and validate operational data first

Reason = Forecasting without clean data leads to unreliable outputs.



HonestAI is a monthly magazine for manufacturers and industrial operators navigating the intersection of growth and artificial intelligence. Every issue is written from the floor up – drawn from real engagements, real problems, and real results inside businesses like yours.

Each issue focuses on a single theme, grounded in real operational case studies and designed to be practical, relevant, and actionable from day one.

Subscribe.

Submit your operation as a case study. Connect with our Editor-in-chief.

HonestAIEngine.com

marketing@honestaiengine.com

+1(470)823-8086

HonestAI

