

## How It Works

*A structured framework for producing inspection-defensible authorization records — built around the question your food & beverage manufacturing inspector will ask, in their language.*

### The question every food & beverage manufacturing authorization eventually faces

*"Who authorized release of this lot, on what basis was the CCP deviation determined not to impact food safety, and why was the decision justified given the documented exception?"*

This is the question an FDA inspector will ask under FSMA — and the equivalent question a third-party certification auditor will ask in different framing. Food safety authorization decisions sit at the intersection of HACCP plans, preventive controls, supplier verification, and recall obligations, where every authorization carries traceability obligations from raw material to finished product.

Most food and beverage quality systems were not built to produce the authorization artifact an inspector or auditor expects to receive. They store HACCP plans, CCP monitoring logs, certificates of analysis, supplier qualification files, and environmental monitoring records — the substrate of the decision. They do not produce a single artifact that captures the authorization logic itself: who decided, on what basis, against what standard, and why.

The Inspection Response Record exists to close that gap. It is a structured artifact — generated in real time, governed by a documented framework, anchored entirely in your inputs — that answers the inspector's question in the form they expect to receive it.

### How the record is produced

<p><b>1</b></p> <p><b>STRUCTURE</b></p> <p>You provide the decision and supporting evidence in</p>	<p><b>2</b></p> <p><b>GENERATE</b></p> <p>The system produces a structured Inspection</p>	<p><b>3</b></p> <p><b>FLAG</b></p> <p>Documentation gaps are identified deterministically</p>	<p><b>4</b></p> <p><b>DEFEND</b></p> <p>You receive a record built to answer the inspector's</p>
--	---	---	--

13 structured fields.	Response Record under strict framework rules.	and surfaced before payment.	question, in their language.
-----------------------	---	------------------------------	------------------------------

**The framework is what makes this defensible.** An AI model on its own would produce variable output. A documented framework with explicit constraints, deterministic flag rules, and a fixed regulatory mapping produces consistent, traceable, inspection-grade records — every time, against every input.

## What the Inspection Response Record Is — and Is Not

The boundaries of this product matter. A Food & Beverage quality leader evaluating an authorization tool needs to know exactly what they are adopting and what responsibility they retain. The framework draws those lines explicitly.

WHAT IT IS	WHAT IT IS NOT
✓ A structured artifact that answers the inspector's exact question — who, when, on what basis.	✗ Regulatory advice or a regulatory opinion.
✓ Generated in real time from your structured inputs about a specific decision.	✗ A consulting engagement, a template fill-in, or a generic compliance product.
✓ Built around the inspector's frame of reference, in inspection-grade language.	✗ A food safety system replacement. It sits above your existing records, not inside them.
✓ Reviewable by a regulatory practitioner before use in any active inspection or enforcement matter.	✗ A certification of compliance or a guarantee that any specific inspection outcome will follow.
✓ Strictly constrained — the system will not introduce evidence, names, dates, or regulations you did not provide.	✗ A free-form AI output. The system is governed by a documented framework that prevents fabrication.

### Where it sits in your food safety system

The Inspection Response Record is positioned above your existing food safety system — not inside it. Your records and supporting documentation continue to operate as they do today. The IRR references those underlying records and produces a structured authorization artifact alongside them.

This means three things for a Food & Beverage quality function:

- Your team continues using the systems they already know. There is no migration, no integration, no rip-and-replace.
- The IRR becomes the artifact you produce alongside any high-stakes authorization decision — finished product release with CCP deviation, raw material acceptance under exception, supplier verification closure, recall initiation, environmental monitoring excursion closure, allergen control determination.

- Your inspector receives the record they were going to ask for, in the structure they were going to ask for it, anchored to the underlying records they already expected to see.

---

*Your food safety system stores the documents. The Inspection Response Record captures the decision.*

---

# How the Record Is Generated

## The framework, the model, and the constraints

The Inspection Response Record is generated in real time by Claude — an AI model produced by Anthropic — operating under a documented framework that constrains its behavior to the structure, language, and logic of an inspection-defensible authorization record.

This combination matters. The model alone would produce variable, free-form output. The framework alone would be a static template. The combination — a model operating under explicit, version-controlled rules — produces a record that is consistent across decisions, traceable to the framework version that governed it, and constrained to the inputs the customer provides.

### What the framework instructs the model to do

1. Answer the inspector's exact question first and directly, using only the customer's structured inputs.
2. Build each section of the record from the corresponding input field — evidence reviewed comes from evidence\_items, justification comes from the justification field, and so on.
3. Apply a fixed regulatory framework based on the customer's industry profile (for Food & Beverage, that framework is 21 CFR 117, 21 CFR 120/123, and applicable third-party certification standards).
4. Run six deterministic gap-detection rules against the inputs and surface every flag explicitly in the output.
5. Maintain a formal, neutral, inspection-grade tone — not marketing language, not generic compliance boilerplate.

### What the framework instructs the model NOT to do

6. Do not introduce evidence, dates, names, titles, or regulations not present in the input.
7. Do not extrapolate beyond what the customer's justification supports.
8. Do not cite regulations from outside the customer's industry framework.
9. Do not paper over missing information with boilerplate or filler.
10. Do not generate a regulatory opinion, a compliance certification, or a determination of inspection outcome.

**This is the trust model.** The framework's defensibility comes from what the system refuses to do as much as from what it produces. A system that fabricates is dangerous. A system that documents what is missing — and refuses to fill the gap with invention — is auditable.

## The Seven Safeguards

The framework imposes seven specific constraints on every record the system generates. These are not aspirational principles — they are mandatory rules embedded directly in the system prompt that governs the model's behavior on every generation.

<b>1</b>	<b>No invented facts</b>	The system is explicitly instructed not to introduce evidence, dates, names, titles, regulations, or rationale not present in your input.
<b>2</b>	<b>Industry-locked regulatory framework</b>	Food & Beverage Manufacturing inputs are cited only against 21 CFR 117, 21 CFR 120/123, and applicable third-party certification standards. The system will not cross-cite into other industry frameworks.
<b>3</b>	<b>Evidence preserved verbatim</b>	The Evidence Reviewed section is built one-to-one from your evidence_items array. It is preserved in the format you provide, not reinterpreted.
<b>4</b>	<b>Justification bounded by your inputs</b>	The Justification Statement cannot exceed what your decision justification supports. The system is instructed not to extrapolate beyond your documented rationale.
<b>5</b>	<b>Gaps surfaced explicitly</b>	When information is missing from a field, the system states it explicitly in Known Limitations and reflects it in the gap count. Silent gaps are not permitted.
<b>6</b>	<b>Deterministic flag logic</b>	Six gap-detection rules run mechanically against your inputs. Flags are not invented or hypothesized — they are applied when specific input conditions are met.
<b>7</b>	<b>Practitioner review available</b>	For records to be used in an active inspection, an open observation, or an enforcement matter, regulatory practitioner sign-off is available through Ostrowski Consulting LLC.

## How a Food & Beverage quality team can verify the framework

Quality leaders evaluating the platform are encouraged to test the framework directly. The recommended evaluation:

11. Run an Inspection Response Record on a deliberately incomplete decision — omit the authorization owner's title, or provide only one piece of evidence, or leave the justification thin. The system should surface every gap explicitly in the flags and in Known Limitations. If it does not, the framework has failed.

12. Run a Record on a complete decision and verify that no fabricated content appears. The Evidence Reviewed section should match your evidence\_items input one-to-one. The regulatory\_alignment section should cite only frameworks within your industry profile. The justification should not exceed what your input supports.
13. Compare the output to a record produced by your own team for the same decision. The structure, language, and depth should be recognizable as inspection-grade. Any divergence is signal — for the framework's calibration or for your team's documentation patterns.

---

*The framework is testable. That is what distinguishes it from a black-box AI tool.*

---

## Common Questions From Quality Leaders

These are the questions a VP Quality / Director of Food Safety at an FDA-regulated food or beverage manufacturer typically asks when evaluating the framework. They are answered here in the same direct form they were asked.

### **Q. Who is generating these records?**

**A.** The structure, framework, and constraints are built by ComplianceWorxs. The text of each record is generated in real time by Claude — Anthropic's AI model — operating under the documented framework rules described in this document. The framework is what governs the output. The model is the engine.

### **Q. Why use AI at all? Why not human authors?**

**A.** Speed and consistency. A practitioner-authored record at this depth would take 5–10 hours and cost \$1,500–\$3,000. The framework-constrained model produces equivalent structure in real time at \$297 per record. For records that require regulatory practitioner sign-off, that review layer is available — but most operational decisions don't need it. The hybrid model scales to the volume of decisions a quality function actually makes.

### **Q. What if the AI hallucinates a regulation or invents a fact?**

**A.** The framework is built specifically to prevent this. The system prompt instructs the model not to introduce regulations, evidence, names, dates, or rationale not present in your inputs. The model writes only what your inputs support. If your inputs are sparse, the record will be sparse and the gap will be flagged — the system will not paper over a gap with fabricated content.

### **Q. How do I know it cited the right regulations?**

**A.** The system applies a fixed regulatory framework based on your industry profile. For Food & Beverage Manufacturing operations, that framework is 21 CFR 117, 21 CFR 120/123, and applicable third-party certification standards. The system will not cite outside this framework. Your regulatory review of the regulatory\_alignment section is part of the intended workflow — the record is a starting point, not a final regulatory determination.

### **Q. Where does the record come from in our food safety system?**

**A.** It sits above your existing food safety system, not inside it. Your records and supporting documentation continue to operate as they do today. The Inspection Response Record references those underlying records — it does not replace them. It is the artifact your team produces alongside the underlying decision, structured to answer the question your inspector will ask.

### **Q. Does this work for both FDA FSMA inspections and SQF/BRC/FSSC 22000 audits?**

**A.** Yes. The framework is designed to produce authorization records that satisfy both FDA inspectors under FSMA and third-party auditors under SQF, BRC, or FSSC 22000. The underlying authorization question — who decided, on what basis, against what standard — is substantively identical across all four. The framework captures the record once; it serves all four audiences.

**Q. Does this work for CCP deviations and recall decisions?**

**A.** Yes. CCP deviation handling, recall initiation determinations, supplier verification closures, and environmental monitoring excursion releases are core authorization decisions the framework is built for. Each is a moment where an inspector or auditor will ask: who authorized this, on what basis, and why was the decision justified given the documented exception?

**Q. Is our data used to train the model?**

**A.** No. Customer inputs and generated records are stored in a secure database under row-level security. They are not used to train the underlying model.

**Q. What if we need a record reviewed by a regulatory practitioner?**

**A.** Practitioner-Reviewed Records are available through a separate review process performed by Ostrowski Consulting LLC. For records being filed in connection with a Form 483, a third-party audit major non-conformity, or a recall determination, practitioner review is recommended.

## FOR FURTHER DILIGENCE

# Two ways to evaluate this further

### 1. Run a record on one of your own decisions

The framework is designed to be tested. Take a recent authorization decision with an exception — a deviation, an OOS investigation, a CCP excursion, a CAPA effectiveness call — and produce an Inspection Response Record on it. The structured output, the gap analysis, and the regulatory alignment will tell you more about the framework in 15 minutes than any document can.

[complianceworxs.com/irr](https://complianceworxs.com/irr) — Inspection Response Record, \$297, generated in real time.

### 2. Request the full procurement-grade framework documentation

The Inspection Response Framework v1.0 is available as a procurement-grade document covering the full system prompt rules, the flag logic with rationale, the change-control governance, and the practitioner review layer. It is intended for regulatory affairs, procurement, and quality management review.

Request from [access@complianceworxs.com](mailto:access@complianceworxs.com).

**For Food & Beverage specifically** — Food and beverage manufacturers operate under FSMA and concurrent third-party certification regimes. The Inspection Response Framework is designed for authorization decisions that must satisfy both FDA inspectors and certification auditors with the same documented logic.

---

**ComplianceWorxs** · The Inspection Response System for FDA-regulated life sciences

[complianceworxs.com](https://complianceworxs.com) · [access@complianceworxs.com](mailto:access@complianceworxs.com)

*Document version 1.0 · Food & Beverage Manufacturing · Effective May 9, 2026 · Subject to change with framework version.  
Each Inspection Response Record carries metadata indicating the framework version under which it was generated.*

*ComplianceWorxs does not make, approve, or recommend regulatory decisions. All determinations remain the sole responsibility of the regulated organization. Use of this system does not constitute legal, regulatory, or compliance counsel.*