

## How It Works

*A structured framework for producing inspection-defensible authorization records — built around the question your 503b outsourcing facility inspector will ask, in their language.*

### The question every 503b outsourcing facility authorization eventually faces

*"Who authorized the release of this CSP batch, on what basis was the environmental excursion determined not to impact sterility, and why was the decision justified at the time?"*

This is the question an FDA investigator will ask when they walk through your 503B facility. With 96% of inspected 503Bs receiving a Form 483, it is the single highest-pressure question in the FDA's inspection universe — and the one most likely to determine whether a finding becomes an observation or escalates to a warning letter.

Most 503B quality systems were not built to answer this question directly. They store batch records, environmental monitoring data, deviation logs, sterility test results, and personnel qualification 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 13 structured fields.</p> | <p><b>2</b></p> <p><b>GENERATE</b></p> <p>The system produces a structured Inspection Response Record under strict framework rules.</p> | <p><b>3</b></p> <p><b>FLAG</b></p> <p>Documentation gaps are identified deterministically and surfaced before payment.</p> | <p><b>4</b></p> <p><b>DEFEND</b></p> <p>You receive a record built to answer the inspector's question, in their language.</p> |
|--|---|--|---|

**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 503B 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 QMS 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 QMS

The Inspection Response Record is positioned above your existing QMS — 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 503B 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 — batch release with deviation, sterility OOS investigation closure, environmental excursion release, media fill failure disposition, CAPA effectiveness for sterility-related findings.
- 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 QMS 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 503B, that framework is 21 CFR 210/211, USP <797>, USP <800>, and FDA 503B guidance).
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> | 503B Outsourcing Facility inputs are cited only against 21 CFR 210/211, USP <797>, USP <800>, and FDA 503B guidance. 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 503B 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 / QA Director at a 503B Outsourcing Facility 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 503B Outsourcing Facility operations, that framework is 21 CFR 210/211, USP <797>, USP <800>, and FDA 503B guidance. 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 QMS?

A. It sits above your existing QMS, 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 sterility OOS investigations and EM excursions?

A. Yes. Sterility OOS dispositions, environmental excursion releases, and media fill failure investigations are exactly the kind of decisions the framework is designed for. Each of these is an authorization

moment with potential patient harm on the table — and each is a question your inspector will ask. The framework captures the data the decision was based on, the regulatory standard applied, the risk evaluation, and the named individual who authorized the disposition.

**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. Records are linked to your email solely for delivery and customer support.

**Q. What about Part 11 — electronic records and signatures?**

**A.** The Inspection Response Record itself is a generated artifact, not a regulated electronic record under 21 CFR Part 11. It is intended to be incorporated into your existing Part 11-compliant quality system through your normal document control process. The record's metadata — session ID, framework version, generation timestamp — supports incorporation into a controlled environment.

**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. Records to be entered into evidence in an enforcement action, used to defend a decision under direct FDA scrutiny, or filed in support of a CAPA closure addressing a prior 483 are recommended for practitioner review.

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 503B specifically** — 503B Outsourcing Facilities operate under the highest-stakes inspection regime in the FDA's universe. The Inspection Response Framework is designed for the decision pressure unique to sterile compounding and the patient safety stakes that follow every batch out the door.

---

**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 · 503B Outsourcing Facility · 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.*