COMPLIANCEWORXS

INSPECTION RESPONSE FRAMEWORK · COSMETICS

How It Works

A structured framework for producing inspection-defensible authorization records — built around the question your cosmetics & personal care inspector will ask, in their language.

The question every cosmetics & personal care authorization eventually faces

"Who authorized this product as adequately substantiated for safety, on what basis was the substantiation determined sufficient, and why was the decision justified at the time?"

This is the question an FDA investigator will ask under the Modernization of Cosmetics Regulation Act (MoCRA). For the first time in the modern era, cosmetics brands face active FDA oversight with mandatory facility registration, product listing, safety substantiation, and adverse event reporting obligations. The authorization questions are new for most brands — but the inspection patterns are not.

Most cosmetics regulatory systems were not built to produce the authorization artifact a MoCRA inspector will expect. They store safety substantiation files, adverse event logs, product listings, facility registration records, and ingredient change controls — 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

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

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 Cosmetics 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 regulatory 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 regulatory system

The Inspection Response Record is positioned above your existing regulatory 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 Cosmetics 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 — safety substantiation determination, adverse event classification (serious vs. non-serious), ingredient or supplier change authorization, responsible person designation, recall initiation, facility registration update.

- 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 regulatory 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 Cosmetics, that framework is FD&C Act §605–§609 (MoCRA framework)).
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.

1	**No invented facts**	The system is explicitly instructed not to introduce evidence, dates, names, titles, regulations, or rationale not present in your input.
2	**Industry-locked regulatory framework**	Cosmetics & Personal Care inputs are cited only against FD&C Act §605–§609 (MoCRA framework). The system will not cross-cite into other industry frameworks.
3	**Evidence preserved verbatim**	The Evidence Reviewed section is built one-to-one from your evidence_items array. It is preserved in the format you provide, not reinterpreted.
4	**Justification bounded by your inputs**	The Justification Statement cannot exceed what your decision justification supports. The system is instructed not to extrapolate beyond your documented rationale.
5	**Gaps surfaced explicitly**	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.
6	**Deterministic flag logic**	Six gap-detection rules run mechanically against your inputs. Flags are not invented or hypothesized — they are applied when specific input conditions are met.
7	**Practitioner review available**	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 Cosmetics 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 Regulatory Affairs / VP Quality at a cosmetics brand operating under MoCRA 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 Cosmetics & Personal Care operations, that framework is FD&C Act §605–§609 (MoCRA framework). 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 regulatory system?

A. It sits above your existing regulatory 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 MoCRA safety substantiation decisions?

A. Yes. Safety substantiation determinations are exactly the kind of decisions the framework is built for. MoCRA §608 requires a brand to establish that products are adequately substantiated for safety — and

an FDA inspector will ask: who reached that determination, on what basis, against what standard, and why was the decision justified? The framework captures that logic in inspection-ready form.

Q. Does this address adverse event classification and reporting under §605?

A. Yes. The classification of an adverse event as serious or non-serious is an authorization decision under §605. Each classification carries reporting obligations and traceability implications. The framework captures the basis for the classification, the evidence reviewed, the named decision-maker, and the rationale — the form an investigator will expect to receive.

Q. MoCRA cGMP rules are still in proposed rulemaking — does the framework reflect that?

A. Yes. The framework references the MoCRA §606 cGMP proposed rule (December 2025) as the cGMP basis when applicable. As FDA finalizes rulemaking, the framework version will be updated to reflect any changes in citation, language, or scope. Each record carries metadata indicating the framework version under which it was generated.

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 an adverse event escalation, a recall initiation, or a substantiation challenge, 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 — 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.

> **For Cosmetics specifically —** Cosmetics brands face MoCRA's regulatory framework still in active rulemaking through 2026 and beyond. The Inspection Response Framework is designed for authorization decisions that must hold up as the regulatory expectations continue to take shape.

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

complianceworxs.com · access@complianceworxs.com

Document version 1.0 · Cosmetics & Personal Care · 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.