

Unified Prompt Contract

The Prompt: four layers. State what the system must do, the boundaries it must respect, and the evidence and output it must produce.

Name	Version	Owner (a person)

FIELD / MEASURE

YOUR ENTRY

Identity & Domain Who the system is, and who it is not; the sources it may use; the determinations it must never make.	
Operational Constraints The never-list, in the policy lane; the instruction-injection stance; the defined paths for refusal, clarification, and escalation.	
Reasoning Protocol The staged, auditable artifacts the model must produce, in order, with the conclusion last.	
Interface Specification The exact output schema - fields, types, allowed values, citation slots, escalation reasons.	

Use concise, auditable evidence and result artifacts, not private reasoning. Validate outputs and enforce consequential permissions in application code.



Unified Prompt Contract

The Record: six fields. Complete this alongside the Prompt. Empty fields are unresolved product decisions.

Name	Version	Owner (a person)

FIELD / MEASURE	YOUR ENTRY
Purpose What this prompt exists to do, and what it must not be used for.	
Inputs Expected fields, optional fields, and behavior when data is missing.	
Outputs Schema, required fields, citation and refusal formats.	
Allowed Actions Tools it may call, tools it may never call, what requires escalation.	
Policy Constraints Tone, safety, compliance-language boundaries - signed by risk.	
Success Criteria The rubric and metrics that define “good” - the hook evaluation attaches to.	

Implementation note: Policy, behavior, and input are product-design categories, not universal API role names. The developer-role examples in Appendix A illustrate an interface that supports that role; map trusted instructions to the provider’s documented hierarchy. Keep user and retrieved-document content separate, test the boundary with hostile inputs, and enforce consequential permissions in application code.

