

GOVERNED TEMPLATE RENDERING REFERENCE: PROMPT ENGINEERING SPEC DEFAULT V2 (DRAFT)

definition af4c3efe...db00

PROMPT ENGINEERING & LLM SPEC · PROJECT ARTIFACT

SPECIRA

Prompt Engineering Spec: Dispatch Modernization

Meridian Field Services: instructional example, not project evidence

DRAFT · WATERMARK POLICY: DRAFT_ONLY

template prompt_engineering_spec v2 · pack: specira_default_delivery

prompts are governed assets · a change without an eval run does not merge

§1 Prompt Inventory

mandatory

1 decision

1 evidence rule

validators: every_prompt_in_registry_with_owner_and_eval_link · every_prompt_has_a_pinned_version_and_changelog_entry · every_prompt_change_has_an_attached_eval_run

One prompt at pilot: a small inventory is an honest inventory. Versions are immutable; rollback is a label move.

Id	Purpose	Target model	Owner	Version · label
PR-01	Rank the eligible technician set for a ready job	<u>the pinned ranking role</u> ^[1]	A. Reyes	v4 · production → v4 (v3 rolled back by label move after its eval i

Anatomy standard: role, instructions, context, examples, output format; PR-01's context block (job attributes plus candidate codes) sits above the instructions per the large-input rule; every candidate row tagged with its source metadata. **Change gate:** a PR-01 diff without an attached eval run fails the pipeline's prompt gate (testing gates^[3]).

§2 Context Strategy

mandatory

1 decision

1 evidence rule

validators: component_budgets_sum_under_window · truncation_ranks_by_relevance_not_recency · window_exclusions_cite_security_or_classification

The smallest high-signal token set, budgeted per component. What never enters the window is a cited design, not a hope.

Component (PR-01)	Budget
System framing	250 tokens
Job attributes	200
Candidate codes (30-candidate ceiling; the eligible set is pre-filtered, <u>BR-01</u> ^[4])	900
Examples	150
Total vs the 1,800 budget (<u>cost architecture</u> ^[5])	1,500

Selection: the eligible set IS the context; no retrieval surface exists for this prompt; past the ceiling, the set truncates by the ranking features (proximity, certification match), never by record recency. **Privacy shape:** technician names, hr_ids, and coordinates NEVER enter the window; anonymized candidate codes only, and the gateway enforces it (SR-014, SR-015^[6], INT-4^[7]): the exclusion is the design.

§3 Injection Defenses

mandatory

1 decision

1 evidence rule

validators: every_untrusted_injection_point_is_delimited_and_typed · defenses_cross_cite_threat_model_rows_both_directions · residual_named_with_adversarial_test_cadence

Defense in depth: the threats are the threat model's, by id; the defenses are this table's; the handshake resolves both ways. Fool-proof prevention does not exist; the residual is named.

Untrusted surface	Defense layers	Threat handshake
Job notes (dispatcher free text, the one untrusted surface reaching PR-01)	Delimited in a tagged data section · typed as an opaque string variable that cannot reshape the template · instruction hierarchy is architectural: the system framing lives at the gateway, above anything the note carries (<u>prompt design</u> ^[5]) · output validation last: schema-constrained ranked ids from the eligible set only; a note that says "rank	<u>TM-04</u> ^[9] (exfiltration) · <u>TM-08</u> ^[9] (over-trust, eliminated); both rows cite these defenses back

Untrusted surface	Defense layers	Threat handshake
-------------------	----------------	------------------

technician X first" cannot mint an id outside the set
([ADR-5](#)^[8], [SR-016](#)^[6])

Residual, named: a crafted note could still bias ranking WITHIN the eligible set; probed by the adversarial suite quarterly ([adversarial set](#)^[2]); accepted rare-low with the override loop as the human backstop ([US-3](#)^[10]).

§4 Evaluation Approach

mandatory

1 decision

1 evidence rule

validators: every_prompt_eval_set_is_cross_referenced_to_test_cases · no_version_live_below_threshold

Every change re-runs the set before merge; shadow deployment before live exposure; a failed run blocks the label move. The history is the lesson.

Design: PR-01's offline set is the golden set ([eval criteria](#)^[1]), threshold rank agreement 0.85; shadow deployment for candidate versions against live pilot traffic, scored async; the user sees production only.

Version	Score	Threshold	Outcome
v2	0.83	0.85	FAILED : label held at v1
v3	0.86	0.85	Shipped; regressed to 0.84 on the refreshed set; label moved back
v4	0.87	0.85	production

§5 Open Questions

mandatory

1 decision

Question	Owner	Answer by	Blocks
Does the quarterly adversarial suite need union-visible reporting? (The note-bias residual touches worker treatment; shared thread ^[6] .)	N. Duval	Sep 30, 2026	The residual's acceptance wording only

Refs References & Package Contents

In the Specira workspace

- | | |
|---------|---|
| specira | [1] Model selection: the pinned ranking role, golden set, eval criteria
app.specira.ai/projects/dispatch-modernization/artifacts/model-selection |
| specira | [2] Test cases: the golden and adversarial sets as registered assets
app.specira.ai/projects/dispatch-modernization/artifacts/test-cases |
| specira | [3] CI/CD pipeline: the prompt change gate app.specira.ai/projects/dispatch-modernization/artifacts/cicd-pipeline#gates |
| specira | [4] Business rules: BR-01 eligibility app.specira.ai/projects/dispatch-modernization/artifacts/business-rules#br-01 |
| specira | [5] AI orchestration blueprint: prompt design, cost architecture
app.specira.ai/projects/dispatch-modernization/artifacts/ai-orchestration-blueprint |
| specira | [6] Security requirements: SR-014, SR-015, SR-016; shared open question
app.specira.ai/projects/dispatch-modernization/artifacts/security-requirements |
| specira | [7] Integration inventory: INT-4 gateway enforcement app.specira.ai/projects/dispatch-modernization/artifacts/integration-inventory#int-4 |
| specira | [8] Decision log: ADR-5 constrained output app.specira.ai/projects/dispatch-modernization/artifacts/decision-log#adr-5 |
| specira | [9] Threat model: TM-04, TM-08 (the threats these defenses answer)
app.specira.ai/projects/dispatch-modernization/artifacts/threat-model |
| specira | [10] User stories: US-3 the human backstop app.specira.ai/projects/dispatch-modernization/artifacts/user-stories#us-3 |

Generated by Specira · template prompt_engineering_spec v2 (draft) · pack specira_default_delivery lineage af4c3efe...db00 · page 1 of 4