

GOVERNED TEMPLATE RENDERING REFERENCE: USER STORIES DEFAULT V2 (DRAFT)

definition 4fbbd943...b192

USER STORIES · REQUIREMENT ARTIFACT

SPECIRA

Story set: Assignment recommendations

Meridian Field Services · Dispatch Modernization: instructional example, not project evidence

DRAFT · WATERMARK POLICY: DRAFT_ONLY

template user_stories v2 · pack: specira_default_delivery depth: standard

§1 Epic Summary & Slice Plan

mandatory 2 decisions 1 evidence rule

The requirement's outcome in one sentence, and the order in which its slices get built, starting with the walking skeleton, the thinnest slice that proves the whole loop works.

Epic. A dispatcher assigns the right technician in under a minute from board suggestions: the story set instantiating the behavior specified in Recommendation rules^[1], feature committed in the project PRD^[2].

DECISION: WALKING SKELETON

US-1 see three suggestions with reasons → **US-2** accept a suggestion → **US-4** assignment recorded and technician notified. The thinnest end-to-end slice: it proves the recommendation loop before any refinement lands. Agreed with A. Reyes (product lead), July 9, 2026.

Slice	Adds	Stories	Maps to
1 (walking skeleton)	The suggestion-to-confirmation loop	US-1 US-2 US-4	Pilot cutover Oct 1 (<u>PRD delivery frame^[2]</u>)
2	Override capture, shortfall handling	US-3 US-7 US-8	Pilot weeks 1 to 2
3 (deferred)	Degraded-feed flagging, corner cases	none	Hub rollout (<u>slice decision^[7]</u>)

§2 Story Register & Acceptance Tests

mandatory

2 decisions

2 evidence rules

validators: ac_rows_carry_concrete_values · preconditions_explicit_even_when_none · boundary_rows_present_for_bounded_inputs

One card per story. Actors are the PRD's personas, never "a user". Each card fences the specified behavior it covers, then proves the fence with acceptance tests: preconditions first (stated even when there are none), concrete test rows against the FRD's scenario outlines (real values, an observable expected result, a path class, a test-case id), then postconditions. Two testers reading the same row must run the same test.

US-1

Dana sees three qualified, nearby suggestions with reasons when a job arrives

actor: Dana (dispatcher) size: 5 depends: none slice 1 vertical slice ✓

covers FR-010, FR-011, FR-013^[1] via outlines AC-010, AC-011^[1]

SCOPE FENCE

Suggestion rendering only. Excludes: acceptance (US-2), override (US-3), shortfall handling (US-7).

PRECONDITIONS

Job JOB-2201 created at the Toronto hub requiring skill "Gas fitting (G2)", routine priority. On-shift technicians positioned: T-014 8 min away, T-022 19 min, T-045 37 min (all G2-certified); T-104 6 min away, on shift, no G2 certification. Dana signed in as dispatcher; Priya signed in read-only.

Row	Path	Test instance	Expected observable result	Test case
AT-1	HAPPY	JOB-2201 arrives on the board	Three suggestions render within 5 seconds, ranked T-014 , T-022 , T-045 , each showing at least two true reasons ("8 min away", "G2 certified")	TC-30
AT-2	BOUNDARY	T-014 repositioned to 9, then exactly 10, then 11 minutes out	Reason chip reads band A ("under 10 min") at 9; band B at 10 and 11. The band edge lands per the FRD's proximity table	TC-30
AT-3	BOUNDARY	T-045 repositioned to 39, then exactly 40, then 41 minutes out	Band C at 39 and 40; at 41 the candidate drops to band D and ranks below every closer qualified candidate	TC-30

Row	Path	Test instance	Expected observable result	Test case
AT-4	NEGATIVE	T-104 is the closest technician but lacks G2	T-104 never appears in the suggestion list; proximity never outranks qualification	TC-30
AT-5	NEGATIVE	Priya (read-only role) opens JOB-2201	Suggestions visible; assignment actions disabled, the FRD's defined permission behavior (ERR-04)	TC-30

POSTCONDITIONS & ENVIRONMENT

Render-only: no assignment state changes; one suggestion-telemetry record per render.
Environment: dispatch board on desktop Chrome and Edge, current versions (PRD experience constraints^[2]).

US-2 Dana accepts a suggestion and the assignment stands

actor: Dana size: 3 depends: US-1 slice 1 vertical slice ✓

covers FR-010 acceptance path via AC-010^[1]

SCOPE FENCE

One-tap acceptance only. Excludes: override (US-3).

PRECONDITIONS

US-1 's state: suggestions rendered for JOB-2201 .

Row	Path	Test instance	Expected observable result	Test case
AT-1	HAPPY	Dana accepts the top suggestion T-014	Exactly one active assignment records T-014 on JOB-2201 ; board shows assigned state	TC-30
AT-2	EDGE	Dana taps accept twice in quick succession	Still exactly one assignment; repeat acceptance changes nothing (idempotent by job id)	TC-30

POSTCONDITIONS

JOB-2201 assigned; one assignment event in the log; notification queued for US-4 's scope.

US-3 Dana overrides with a reason code and her choice wins

actor: Dana size: 3 depends: US-1 slice 2 vertical slice ✓ agent-flagged

covers FR-012 via outline AC-012^[1]

SCOPE FENCE

Override capture with a stored reason. *Excludes:* reason-code configurability, open question **0Q-2** , A. Reyes, Jul 25 ([FRD open questions^{\[6\]}](#)).

PRECONDITIONS

Job **JOB-2214** open at the Toronto hub with three suggestions rendered; Dana signed in as dispatcher; **T-088** on shift but not among the suggestions; **T-104** on shift, missing the required gas certification.

Row	Path	Test instance	Expected observable result	Test case
AT-1	HAPPY	Dana assigns T-088 with reason code RC-03 "customer request"	Assignment records T-088 with RC-03 ; T-088 notified; override logged in suggestion telemetry	TC-31
AT-2	NEGATIVE	Dana confirms an override with no reason code selected	Assignment does not complete; the board states a reason is required; no assignment event is recorded	TC-31
AT-3	BOUNDARY	Optional override note at exactly 250 characters, then 251	250 records in full; 251 is rejected with the limit named, with no silent truncation	TC-31
AT-4	EDGE	Dana overrides to T-104 , who lacks the required gas certification	The board warns, requires explicit confirmation, and records the skill-mismatch flag: the FRD's defined behavior, and her choice still wins (FR-012^[1])	TC-31

POSTCONDITIONS & ENVIRONMENT

JOB-2214 assigned; the event log holds one override entry with its reason code; suggestion telemetry updated. Environment: board on desktop Chrome and Edge current; the technician-facing confirmation rides **US-4** 's mobile scope.

US-4 **Marcus receives and confirms his assignment on his phone**

actor: Marcus (technician) size: 5 depends: **US-2** slice 1 vertical slice ✓

SCOPE FENCE

Delivery and one-tap confirmation. *Excludes:* no-show flagging (later requirement).

PRECONDITIONS

An assignment event for **T-088** (Marcus) exists; his day view is signed in on a mobile device.

Row	Path	Test instance	Expected observable result	Test case
AT-1	HAPPY	Assignment of JOB-2214 to Marcus records	The job appears on his day view within 30 seconds with address and skill shown	TC-311
AT-2	HAPPY	Marcus taps confirm	Confirmation is visible to Dana on the board within 5 seconds	TC-311

POSTCONDITIONS & ENVIRONMENT

Confirmation event logged. Environment: day view on iOS Safari and Android Chrome, current versions ([PRD experience constraints^{\[2\]}](#)).

SP-1 Spike: historical drive-time data quality

time-box: 2 days outcome: go/no-go note on band accuracy feeds [distance-proxy assumption^{\[8\]}](#)

DECISION: SPLIT APPLIED & ROW DEFERRALS

Suggestion generation and acceptance were one 13-point story; split by workflow step into **US-1** and **US-2** ; the register's split ceiling is 8 (§4). Corner-case rows (feed loss mid-override) are deferred to slice 3 by name ([slice decision^{\[7\]}](#)). A deferral is a decision; silence is not.

§3 Acceptance Tests & the FRD Outline Bridge

mandatory 1 decision 1 evidence rule

validators: no_gherkin_syntax_in_story_ac · coverage_map_resolves_both_ways · ac_rows_link_bidirectionally_to_test_cases

Three layers, captured once each. The FRD owns every rule as an abstract scenario outline, the only Given/When/Then anywhere. Stories own concrete rows against those outlines: the tables in §2. Execution evidence (§5) proves the rows ran. A rule discovered during story writing goes up to the FRD as a change request; a new data instance stays here as a row.

FRD outline / behavior	Story rows	Test cases	Status
AC-010: ranked suggestions & acceptance	US-1 AT-1..3 · US-2 AT-1..2	TC-301..303, 306..307	covered

FRD outline / behavior	Story rows	Test cases	Status
AC-011: visible reasons	US-1 AT-1	TC-301	covered
AC-012: override with stored reason	US-3 AT-1..4	TC-311..314	covered
FR-013: unqualified never proposed	US-1 AT-4	TC-304	covered
ERR-01 / ERR-06: late, shortfall	US-7 AT-1..3	TC-321..323	covered
ERR-02 / ERR-03: no match, bad skill	US-8 AT-1..2	TC-331..332	covered
ERR-04: permission denial	US-1 AT-5	TC-305	carried
ERR-05: legacy screen after cutover	deferred to the cutover slice: decision ^[9]		deferred

EVIDENCE: THE BOUNDARY HELD

One missing rule surfaced during example mapping (a double-acceptance race) and was filed as an FRD change request rather than written into US-2 . The FRD stays the single source of rules; the story kept only the concrete instance (US-2 AT-2) once the outline landed.

§4 Sizing, Example Mapping & Readiness Gates

mandatory

2 decisions

1 evidence rule

validators: example_mapping_precedes_ready · ready_gate_scoped_to_dev_entry · spikes_timeboxed_never_pointed

Honest sizes, the mapping session that authored the rows, and two distinct gates. Ready protects entry to development; it is deliberately not a full-specification stage gate, so refinement flow never blocks. Done protects closure (§5 holds the evidence). The lifecycle renders as a state diagram because this story set customizes the gates.

Scale: modified Fibonacci; split ceiling 8 for this set (agreed with the team, July 9). Spikes are time-boxed with a stated knowledge outcome, never pointed: SP-1 is two days, its


output a go/no-go note.

EVIDENCE: EXAMPLE MAPPING, JULY 9, 2026

Three Amigos per story, 25 minutes each: A. Reyes (product), E. Sandoval (engineering), K. Yamada (QA lead; QA in the room is the point). Four rules mapped to FRD links, eleven examples became §2's acceptance-test rows, two questions moved to §8 with owners the same day. **US-6** overran its timebox; split before Ready, recorded as the signal it is.

Ready (entry to dev): example-mapped · rows concrete · size set · dependencies stated · FRD links resolve. **Done:** the project standard ([Definition of Done^{\[5\]}](#)) plus one story-set addition: override-reason telemetry visible on the pilot dashboard. Closure itself is governed by §5's evidence and sign-off.

SAMPLE
seeded demo data · specira.ai



Story lifecycle: state diagram rendered through Specira's diagram service. Fallback: Draft → Ready (DoR gate: example-mapped, concrete rows, size, dependencies, FR links) → In progress → In review (verification passes; findings loop back) → Done (DoD: project standard + story-set additions; acceptance evidence recorded).

§5 Acceptance & Sign-off

mandatory

2 decisions

1 evidence rule

validators: signoff_evidence_recorded_at_closure · defects_carry_disposition_at_acceptance

A story is accepted on evidence, not on a nod. Three distinct confirmations: the product owner accepts business fit, the domain reviewer confirms the workflow reads true to practice, QA confirms every row executed with results recorded. Failed acceptance returns through the same Ready gate, never forced, and evidence is never edited after sign-off.

Evidence record: US-3 (accepted September 18, 2026)

Test case	Row	Result	Environment / build	Tester	Date
TC-311	AT-1	pass	pilot · build 2026-09-17	K. Yamada	Sep 18
TC-312	AT-2	pass	pilot · build 2026-09-17	K. Yamada	Sep 18
TC-313	AT-3	fail → pass (DF-112)	pilot · builds 2026-09-16 / 09-17	K. Yamada	Sep 16 / 18
TC-314	AT-4	pass	pilot · build 2026-09-17	K. Yamada	Sep 18

Defect	Found by	What	Disposition
DF-112	TC-313	Override note truncated at 249 characters	Fix before acceptance: fixed in build 2026-09-17, re-run passed
DF-114	TC-311	Reason-code chip overlaps the technician name at narrow widths (cosmetic)	Defer: owner D. Kaur (design), due October 9, 2026

SIGN-OFF: ACCEPTED

Business fit: A. Reyes (product lead) · workflow true to practice: M. Chen (dispatch lead) · coverage and execution: K. Yamada (QA lead). September 18, 2026. Full run export attached (US-3 acceptance run^[4]); test-case links resolve in the test-case register^[10].

KICKBACK: THE GATE WORKS BOTH WAYS

US-7 failed its first acceptance on TC-322 (shortfall notice missing the count) and returned to the backlog through the same Ready gate; the failed run's evidence is retained unedited beside the passing one.

§6 Agent Implementation Handoff

conditional · active: stories may be agent-implemented

1 decision

1 evidence rule

validator: agent_handoff_complete_when_flagged

A thin story makes a coding agent guess, and agents guess plausibly-wrong. Every agent-flagged story carries pointers, boundaries, runnable verification, and one pattern to imitate. The test: an agent told to ask clarifying questions first should have none these fields don't already answer.

story: US-3: override with reason code **pointers:** board assignment module · override dialog component · assignment event recorder **exemplar:** the acceptance path shipped in US-2; imitate its event pattern **boundaries:** always: record the reason code atomically with the assignment ask first: any new dependency or schema change never: touch suggestion ranking **verify:** run the board suite's override cases; they implement TC-311..314, so the agent's checks and QA's checks agree · manual mirror of row AT-2: an override without a reason cannot complete **pitfall:** acceptance is idempotent by job id; do not introduce a second write path

DECISION: AGENT READINESS

US-3 is agent-implementable as specified. US-1 stays human-paired: it establishes the board's suggestion rendering patterns that later stories imitate, and the exemplar has to exist before agents can follow it.

§7 Edge & Negative-Path Stories

mandatory

1 decision

2 evidence rules

validators: error_catalog_fully_mapped · negative_rows_present_or_explicitly_deferred

The FRD's error catalog, story by story: covered, carried inside a feature story's rows, or deferred with a slice. Every bounded input gets its boundary rows somewhere in the set. No failure behavior is invented at story level: rows instantiate the FRD's defined failure semantics with real triggering values.

US-7

Dana sees an honest shortfall notice when fewer than three technicians qualify

actor: Dana size: 3 slice 2 covers ERR-06, ERR-01 render path^[1]

PRECONDITIONS

Jobs staged so exactly two, one, and zero technicians qualify (certification and shift data set per row).

Row	Path	Test instance	Expected observable result	Test case
AT-1	NEGATIVE	Two of three qualify for JOB-2230	The FRD's defined shortfall notice renders with the count named: no invented copy, no silent padding with unqualified names	TC-32
AT-2	NEGATIVE	One qualifies	Notice names the single candidate and the shortfall count	TC-32
AT-3	NEGATIVE	Zero qualify	Empty-state per the FRD: the missing skill named, no suggestions fabricated	TC-32

US-8 Intake rejects an unknown skill with the offending value named

actor: Dana size: 2 slice 2 covers ERR-02, ERR-03^[1]

ROWS

AT-1: job submitted with skill "Gas fitting (G4)" (not in the taxonomy) → rejected naming the value (TC-331) · AT-2: empty skill field → the FRD's required-field behavior (TC-332).

Boundary coverage & path classes across the set

Bounded input	Boundary rows	Lives in
Proximity bands (FR-011 ^[1] : A<10, B 10 to 20, C 20 to 40, D>40 min)	9 / 10 / 11 · 39 / 40 / 41	US-1 AT-2, AT-3
Override note length (250-character maximum)	250 / 251	US-3 AT-3

Story	Happy	Negative	Boundary	Edge	Corner
US-1	●	●	●	none	deferred to slice 3 by name (feed loss mid-action): decision ^[7]
US-2	●	none	none	●	
US-3	●	●	●	●	
US-7 / US-8	none	●	none	●	

§8 Open Questions

mandatory

1 decision

Owned and dated, with what each blocks; an unowned question is not tracked. Questions raised in example mapping land here the same day.

Question	Owner	Answer by	Blocks
OQ-2 Are override reason codes tenant-configurable or fixed? RAISED IN THE JULY 9 EXAMPLE-MAPPING SESSION	A. Reyes (product lead)	Jul 25, 2026	US-3 's reason-code list only (<u>FRD^[6]</u>)

No question blocks the walking skeleton.

Omission note: how the agent-handoff section renders when a team is all-human: "§6 Agent Implementation Handoff omitted: implementation is exclusively human this cycle. Rationale recorded in adaptation event #2."

Refs

References & Package Contents

Package items travel in the same export bundle; workspace items link into Specira where the governed record lives.

In this export package

PACKAGE	[1] Functional spec: Recommendation rules (FR, ERR, and outline identifiers) ./frd-recommendation-rules.docx
PACKAGE	[2] Product Requirements Document: Dispatch Modernization ./prd-dispatch-modernization.docx
PACKAGE	[3] Story-map board export (backbone and slices) ./story-map-recommendations.html
PACKAGE	[4] US-3 acceptance run export (results, defects, sign-off) ./acceptance-runs/us-3-run-2026-09-18.html

In the Specira workspace

SPECIRA	[5] Knowledge base: Definition of Done (project standard) app.specira.ai/kb/entries/definition-of-done
SPECIRA	[6] FRD open question: override reason codes app.specira.ai/projects/dispatch-modernization/artifacts/frd-recommendation-rules#open-questions
SPECIRA	[7] Decision: slice plan accepted (A. Reyes, July 9, 2026) app.specira.ai/projects/dispatch-modernization/decisions/95
SPECIRA	[8] Assumption: distance proxy, shadow trial RQ-041 app.specira.ai/projects/dispatch-modernization/requirements/RQ-041

- SPECIRA

[9] Decision: pilot cutover plan (ERR-05 deferral)

app.specira.ai/projects/dispatch-modernization/decisions/93
- SPECIRA

[10] Test-case register: TC identifiers resolve here, linked back to story rows

app.specira.ai/projects/dispatch-modernization/test-cases

Generated by Specira · template user_stories v2 (draft) · pack specira_default_delivery
lineage 4fbbd943...b192 · page 1 of 11

SAMPLE
seeded demo data · specira.ai