

GOVERNED TEMPLATE RENDERING REFERENCE: JOURNEY MAPS DEFAULT V2 (DRAFT)

definition 11f1cf84...67d1

JOURNEY MAPS · PROJECT ARTIFACT

SPECIRA

Journey Maps: Dispatch Modernization

Meridian Field Services: instructional example, not project evidence

DRAFT · WATERMARK POLICY: DRAFT_ONLY

template journey_maps v2 · pack: specira_default_delivery

one persona · one scenario · one goal per map

§1 Core Journey Pairs: Dana (JM-01/02 linear · JM-06/07 non-linear)

mandatory

3 decisions

1 evidence rule

validators: future_state_declares_current_pair · delta_rows_per_shared_stage · per_stage_metrics_same_definition_baseline_target_owner · non_linear_scenarios_declare_branch_loop_or_handoff_mode

One spine, two maps. The current state is evidence; the future state is measurable requirements. Per stage: what changes, the capability that delivers it (a label, not a blueprint), and the expected shift. Every delta traces to a pain; every target traces to a committed number. A to-be map that is the as-is with nicer emotions fails review.

Pair scope: PER-001 Dana^[1] · a priority job arrives mid-morning · assigned inside a minute. **JM-01** current, research-first (six interviews, two floor observations, June 2026^[6]); **JM-02** future, pair **JM-01** , same four-stage spine, horizon: pilot through hub rollout, March 2027. Owner: D. Kaur; JM-01 validated July 10, 2026.

JM-01: current state (evidence-based)

	Intake call	Find technician	Assign & confirm	Paper trail
ACTIONS	Capture fault; promise a window	Radio poll; whiteboard check	Radio confirm; rekey into the scheduler	Log the choice from memory if challenged

	Intake call	Find technician	Assign & confirm	Paper trail
THOUGHTS	"Which crew is even close?"	"I'm guessing on positions"	"Did the rekey take?"	"Hope nobody asks why"
EMOTION (BASIS)	Strained (observed, both observations)	Low point (observed)	Relief with residue (inferred, basis stated)	Exposed (interviews 4, 6)
TOUCHPOINTS	Phone · legacy scheduler	Radio · whiteboard	Radio · scheduler	Paper log
MOMENT OF TRUTH	none	Customer on hold while Dana guesses, where windows are lost and the idle cost accrues (<u>problem evidence^[2]</u>)	none	none

JM-02: future state, same spine, with deltas and targets


	Intake call	Find technician	Assign & confirm	Paper trail
WHAT DANA DOES	The job lands on the board as the call ends	Reads three suggestions with visible reasons	One tap accepts, or a coded override	Nothing: the log wrote itself
DELTA: WHAT CHANGES · CAPABILITY	No rekeying · board intake (<u>feature inventory^[3]</u>); traces to pain 1	Guessing becomes verifying · recommendations with reasons (<u>FR-010</u> , <u>FR-011^[4]</u>); traces to pain 2	Two systems become one action · acceptance + override capture (<u>US-2</u> , <u>US-3^[5]</u>)	Memory becomes record · override log (<u>US-3^[5]</u>); traces to pain 3
EXPECTED SHIFT	Strained → steady	The low point becomes the strong point: the customer stays on the line; the moment of truth inverts	Relief without residue	Exposed → covered
METRIC: ONE DEFINITION, BASELINE →	Job-visible latency: none	Time-to-assign: 8 to 15 min observed →	Override rate: n/a → tracked,	Idle-time share: BRD

	Intake call	Find technician	Assign & confirm	Paper trail
TARGET (OWNER P. NOVAK)	(phone-relayed) → ≤5 s p95 (quality scenario ^[7])	<1 min (committed ^[3])	stable band (telemetry ^[5])	baseline → ≤10% by Mar 31, 2027 (the committed number ^[2])


ASSUMPTIONS: OWNED, SCHEDULED, INSIDE THE HORIZON

Feed availability holds at 99.5% (owner S. Grewal, validated against INT-1 reporting, November 2026). Dispatchers adopt suggestions rather than bypassing (owner M. Chen, validated by override-rate telemetry, pilot weeks 1 to 4). Union accepts override-telemetry scope (owner N. Duval, August 8 answer). Constraint carried: propose-only posture (ADR-3^[7]). Anything past March 2027 is out of scope, not aspirationally included.

The improvement thesis in one picture: both arcs, one axis



JM-01 current: the dip at "find technician" is the moment of truth failing.



JM-02 future, same stages: the dip becomes the plateau; the targets in the matrix say by how much and by when. Both arcs render from committed source.

The same laws on a scenario that is not a straight line: JM-06 · JM-07

Ordered is not linear. A branch, a rework loop and a cross-persona handoff all fit the pair discipline; each declares the mode it chose rather than being flattened into a straight line. The simple pair above and this one are the two ends of the range an intaker has to cover.

Pair scope: [PER-001 Dana^{\[1\]}](#) · a callback raises a job's priority while a technician is already rolling to another · re-dispatch without breaking the committed window on *either* job.

JM-06 current, research-first (four interviews plus one floor observation, June 2026 ^[6]);

JM-07 future, pair **JM-06**, same five-stage spine, horizon: March 2027. Owner: D. Kaur.

COMPLEXITY TRIGGERS: WHAT FIRED, AND THE MODE EACH CHOSE

Branch at "Choose response": three outcomes that all rejoin the same goal, so they ride **one spine as per-outcome rows**, not three maps. **Handoff**, not a fourth row: the split-crew outcome needs another persona's decision, and work that crosses a persona is a link (**JM-09**), never a row. **Loop** at "Secure acceptance": a declined reassignment returns to "Choose response" as a **return edge** carrying its frequency and its exit, not a stage. Because the loop is an edge, the spine stays five stages on both maps, which is what lets the two arcs share one axis.

	Callback lands	Assess in-flight state	Choose response	Secure acceptance	Settle both jobs
ACTIONS (JM-06)	Take the callback; re-rank by hand	Radio the technician; guess an ETA	Weigh three options on the whiteboard	Radio the reassignment; absorb a decline	Rekey both jobs; call the bumped customer
THOUGHTS	"What do I have to break?"	"Where is he actually?"	"Whichever I pick, someone waits"	"Now I start over"	"Hope nobody asks why"
EMOTION (BASIS)	Strained (observed)	Anxious (observed)	Torn (interviews 1, 4)	Low point: the decline loop bites (observed)	Exposed (interviews 4, 6)
TOUCHPOINTS	Phone · legacy scheduler	Radio	Whiteboard	Radio	Scheduler · phone · paper log
DELTA: WHAT CHANGES · CAPABILITY (JM-07)	No reconstruction from memory · board intake (<u>feature inventory</u> ^[3]); traces to pain 1	Guessing becomes reading · position feed (<u>FR-010</u> ^[4]); traces to pain 2	The branch stays; guessing between branches becomes comparing them · branch cost comparison (<u>FR-011</u> ^[4]); traces to pain 2	The loop shrinks; a decline now carries a reason that feeds the next suggestion · acceptance + reason capture (<u>US-2</u> ^[5])	Memory becomes record · override log + customer notice (<u>US-3</u> ^[5]); traces to pain 3


	Callback lands	Assess in-flight state	Choose response	Secure acceptance	Settle both jobs
METRIC: BASELINE → TARGET (OWNER P. NOVAK)	Time-to-redispatch: 12 to 25 min observed → <3 min (<u>experience constraints^[3]</u>)		Decline-and-retry rate: ~1 in 3 (interviews 1, 4) → ≤1 in 10 ; the loop is waste, so it is measured as waste		Bumped-window rate: BRD problem evidence → ≤2% by 31 Mar 2027 (<u>success criteria^[2]</u>)

The branching stage: a row per outcome, each with its own moment of truth


Outcome	Condition that selects it	Moment of truth (differs per outcome)	Delta in JM-07
Reassign	The rolling job has not started	The first customer's window breaks, and Dana makes that call herself	The board prices the broken window before she chooses
Finish first	The rolling job is nearly done	The priority caller is told "within the hour"; the escalation bought nothing	Remaining-time is on screen, so "nearly done" stops being a guess
Split the crew	A two-technician job	Both clocks run while an exception waits for approval	Approval is requested in-flow with the cost attached → JM-09 (handoff, not a branch)

ASSUMPTION: OWNED, SCHEDULED

Technicians give a reason rather than a bare decline (owner M. Chen, validated by reason-completeness telemetry, pilot weeks 1 to 4). If it fails, the loop delta is the first thing to lose.



JM-06 current: the low point is "secure acceptance", where the decline loop bites. The arc renders the **dominant path**: mermaid's journey notation has no way to say "one of these three", so the branch stays in the matrix where it is legible.



JM-07 future, same five stages: the loop's low point lifts. Same axis as JM-06 because the spine never changed: the branch is rows, the loop is an edge.

§2 Journey Portfolio & Handoffs

mandatory

1 decision

1 evidence rule

validators: portfolio_covers_every_primary_core_scenario ·
handoff_links_carry_target_map_and_stage_both_ways ·
maps_carry_owner_status_review_cadence

Maps without governance are cemeteries. The portfolio says which maps exist and why, who owns each, and how personas' journeys hand off to each other, so improving a shared touchpoint

shows its cascade.

Map	Persona × scenario	State	Pair	Owner	Status	Review
JM-01	Dana × priority job	current	JM-02	D. Kaur	validated	quarterly
JM-02	Dana × priority job	future	JM-01	D. Kaur	targets committed	quarterly; re-measure at cutover
JM-03	Marcus × assignment day	current	JM-04	D. Kaur	validated (4 ride-alongs)	quarterly
JM-04	Marcus × assignment day	future	JM-03	D. Kaur	targets committed	quarterly
JM-05	Dana × feed-loss recovery	future	none (scoping note: the current analog, a radio blackout, differs structurally)	D. Kaur	assumptions open	per drill
JM-06	Dana × escalation mid-flight	current	JM-07	D. Kaur	validated	quarterly
JM-07	Dana × escalation mid-flight	future	JM-06	D. Kaur	targets committed	quarterly; re-measure at cutover
JM-08	Priya × exception approval	current	JM-09	D. Kaur	validated (2 interviews)	twice yearly
JM-09	Priya × exception approval	future	JM-08	D. Kaur	targets committed	twice yearly

DECISION: COVERAGE, THE RULE APPLIED

Both primary personas' core scenarios are paired, no exceptions. **The coverage rule is scoped to a scenario, not to a persona.** Priya (secondary) rides **JM-01/02** as a read-only layer: in the priority-job scenario her stages do not diverge, rationale recorded. In the

escalation scenario they do (a split-crew exception hands her a minute-scale, time-boxed approval her routine review never has), so she earns **JM-08/09** for that scenario and that one alone ([PER-002 exception duty^{\[1\]}](#)). Diverging once does not promote a secondary persona everywhere. **JM-03/04** follow the same pair anatomy, compressed here. Signature delta: the paper day sheet becomes one-tap confirmation; metric (time-to-confirm): baseline end-of-day paper → **≤5 min p90** ([US-4 rows^{\[5\]}](#)). **JM-08/09** the same way. Signature delta: the approval request arrives by phone with no context and she approves blind or stalls; in JM-09 it arrives with the branch's cost attached and resolves in one tap, logged; metric (time-to-approve): 6 to 20 min phone tag (interviews 7, 8) → **≤2 min p90**, owner P. Novak.

HANDOFF LINKS: BOTH DIRECTIONS

JM-02 "Assign & confirm" **triggers** **JM-04** "Receive next job", and JM-04 declares the trigger back. The shared touchpoint is the assignment event ([Assignment^{\[8\]}](#)): an improvement there cascades across both maps, visibly. **JM-07** "Secure acceptance" **triggers** that same **JM-04** stage, and JM-04 declares **both** triggers back (the routine assignment from **JM-02** and the escalated one from **JM-07**) on that same shared touchpoint. The escalation reuses the technician's existing map rather than minting a second one because his stages do not change when the job reaching him was escalated: only the trigger differs, and a differing trigger is a link, not a new map. **JM-07** "Choose response" **triggers** **JM-09** "Approve the exception" on the split-crew outcome only, and JM-09 declares that trigger back; the shared touchpoint is the exception request ([Exception^{\[8\]}](#)).

Wrong-tool omissions, recorded: Priya's weekly exception review has no inherent order (she enters it by exception type, by age, or wherever she left off) → the screen inventory owns it. Her *in-flight* approval is a different scenario and maps as **JM-08/09**: one persona can own both an ordered scenario and an orderless one, so the omission is recorded per scenario, never per persona. Warehouse flows are out of scope ([PRD^{\[3\]}](#)). **Portfolio review:** with the monthly stakeholder refresh; post-cutover, JM-01/03/06/08 archive as baselines and the pairs re-measure apples to apples.

§3 Onboarding Journey

conditional · omitted

validator: onboarding_conditional_on_first_run_trigger

Omitted with rationale: dispatchers are trained staff on a managed rollout; no self-serve first-run exists. The cutover training plan owns the first-session experience ([Pilot cutover plan^{\[9\]}](#)); activation is the pilot readiness gate, not a product journey. Recorded in the portfolio and in adaptation event #1.

§4 Error-Recovery Journey: feed loss (JM-05)

conditional · engaged: high-stakes failure

1 decision

1 evidence rule

validator: error_recovery_conditional_on_stakes_trigger

The highest-stakes failure carries the map. The FRD's defined behavior is referenced, never re-specified; this map narrates how the failure feels and where trust is won or lost.

	Notice	Continue	Recover
WHAT HAPPENS	The board flags stale positions honestly	Assignment flow proceeds without proximity (<u>degradation^[4]</u>)	Positions return; the flag clears
EMOTION (BASIS)	Alert, not alarmed (interview 2: "just tell me it's old")	Steady if informed  inferred; validation at the first feed drill	Restored
MOMENT OF TRUTH	The honest flag: trust is won by the flag, lost by silent wrongness	none	none
RECOVERY METRIC	Zero blocked assignments during loss; baseline n/a (new capability), target 0 , owner K. Yamada (<u>feed-loss scenario^[7]</u>)		

Opportunities: the staleness flag's copy → open question below · recovery telemetry on the pilot dashboard → US-3 Done addition^[5].

§5 Open Questions

mandatory

1 decision

Question	Owner	Answer by	Blocks
Validate the "steady if informed" inference at the pilot's first feed drill	D. Kaur	Oct 16, 2026	The staleness-flag copy only (<u>ERR catalog copy^[4]</u>)
JM-02 assumption checks: feed availability, adoption, union scope	S. Grewal · M. Chen · N. Duval	Nov 2026 · pilot wks 1 to 4 · Aug 8	None blocks design use of the pair

Refs

References & Package Contents

In this export package

- package [2] Business Requirements Document: problem evidence, success criteria ./brd-dispatch-modernization.docx
- package [3] Product Requirements Document: feature inventory, solution overview ./prd-dispatch-modernization.docx
- package [4] Functional spec: degradation, error catalog ./frd-recommendation-rules.docx

In the Specira workspace

- specira [1] Persona profiles: PER-001 Dana (attributes live there, referenced by id)
app.specira.ai/projects/dispatch-modernization/artifacts/persona-profiles
- specira [5] Story set: US-3, pilot dashboard Done addition app.specira.ai/projects/dispatch-modernization/artifacts/user-stories
- specira [6] Discovery sessions 2 and 3: the research base app.specira.ai/projects/dispatch-modernization/discovery/sessions
- specira [7] Architecture: quality scenarios (latency, feed-loss), ADR-3
app.specira.ai/projects/dispatch-modernization/artifacts/architecture
- specira [8] Data model: the Assignment event and the Exception request (the shared touchpoints)
app.specira.ai/projects/dispatch-modernization/artifacts/data-model
- specira [9] Decision: pilot cutover plan (training owns first-run) app.specira.ai/projects/dispatch-modernization/decisions/93