

GOVERNED TEMPLATE RENDERING REFERENCE: PRD DEFAULT V2 (DRAFT)

definition a4c029aa...d46e

PRODUCT REQUIREMENTS DOCUMENT · PROJECT ARTIFACT

SPECIRA

Dispatch Modernization

Meridian Field Services: instructional example, not project evidence

DRAFT · WATERMARK POLICY: DRAFT_ONLY

template prd v2 · pack: specira_default_delivery depth: standard

§1 Executive Summary

mandatory

written last · introduces nothing new

One paragraph a sponsor can repeat in an elevator: the problem, what we are doing about it, and the single number (with its date) that tells us it worked.

Meridian's dispatchers coordinate 400 field technicians using a scheduling screen built in 2009 and a lot of phone calls. The result: technicians sit idle for 18% of their paid shift, roughly **\$2.1 million a year** in hours Meridian pays for but cannot bill, of which about **\$700 thousand is assignment delay**, the share this project removes (idle-cost analysis^[10]). Dispatch Modernization replaces manual assignment hunting with a recommendation built from each technician's skills and current location, presented inside one live dispatch board. The Toronto pilot cuts over on **October 1, 2026**; success means idle time at **10% or lower by March 31, 2027**, two quarters after pilot cutover, measured with the vehicle-tracking data Meridian already collects today.

§2 Problem & Product Vision

mandatory

3 decisions

2 evidence rules

Why this project exists, in the customer's own words, and what staying with the status quo costs. No solution language is allowed in this section; that discipline is enforced by the template.

"I juggle three screens and a phone headset to find who's free, who's qualified, and who's closest; by the time I've decided, the job's aged twenty minutes."

Dispatcher interview, June 12, 2026^[6]

Meridian's technicians are idle 18% of the time they are on shift, against a 9% industry benchmark. That gap costs an estimated \$2.1 million a year in unbilled hours, and slow dispatch is the most-cited reason enterprise customers give when they leave (exit-interview themes, Q2 2026^[7]).

Why 10% and not the 9% benchmark: the benchmark assumes route optimization across multi-day schedules, which is explicitly out of this project's scope (§5). This project removes the assignment-delay share; the remaining point belongs to the Routing initiative.

Why now. Meridian's two largest customer contracts re-tender in **January 2027**. Both customers named response time in their last business reviews. A board live from October 1 gives a full quarter of measured improvement before those tenders open.

The world once solved. A dispatcher opens one board, sees every job and every technician in one place, and gets a sensible suggestion the moment a job arrives. Assigning work takes seconds of confident judgement instead of minutes of detective work.

DECISION: PROBLEM FRAMING

The one problem this project commits to solving: **assignment latency and mismatch in dispatch**. Where the dispatcher's needs and the technician's preferences conflict, the dispatcher's speed wins. At launch the technician is protected by the propose-only posture: a person confirms every assignment (§10); fairness reporting ships in the second slice, January 2027 (§7).

EVIDENCE

Idle-time baseline from the vehicle-tracking feed, weekly exports May to June 2026 · delay share from the idle-cost analysis^[10] · churn attribution from three enterprise exit interviews (knowledge base^[7]).

§3 Target Users & Jobs-to-be-Done

mandatory

2 decisions

1 evidence rule

Who this is for, described as real people with a job they are trying to get done. Every other section refers to these three by name.

Dana · primary

Dispatcher, Toronto hub

"Get the right technician moving in under a minute, without second-guessing myself."

Assigns 300+ jobs a day. Today's workaround: three screens, a paper skills cheat-sheet, and calling technicians to ask where they are.

Priya

Operations manager, Eastern region

"Show me today's exceptions before the customer phones me about them."

Today's workaround: an end-of-day spreadsheet assembled from dispatcher memory.

Marcus

Field technician, HVAC crew

"Give me a full, sensible day of work, not a zigzag across the city with two reassignments before lunch."

Today's workaround: negotiating jobs by text message.

DECISION: PERSONA PRIORITY

Dana is the primary persona: dispatch speed is the bottleneck the business case rests on. Deliberately excluded from this project: Meridian's end customers (no customer-facing surface ships; see [Scope, §5](#)) and the finance team (billing integration waits for the finance system replacement, Q3 2027).

EVIDENCE

All three personas come from discovery sessions with eleven staff at the Toronto and Halifax hubs, May to June 2026 ([persona records^{\[8\]}](#)), the discovery sample, distinct from the rollout sequence in §5. None are invented.

§4 Success Metrics & Instrumentation

mandatory

3 decisions

2 evidence rules

The scoreboard, agreed before anything is built, including how each baseline gets captured and what, precisely, makes a guardrail trip. "Leading" means an early signal we can act on within weeks; "lagging" means the outcome that confirms it later.

What we measure	In plain terms	Baseline	Target	Target date
Technician idle rate NORTH STAR	Share of a paid shift with no job assigned	18%	≤ 10%	Mar 31, 2027 4-1
Assignment speed GUARDRAIL · EXPERIENCE	Time from a job arriving to an accepted assignment, for three of four jobs, the same measure the business case uses (BRD ^[1])	Sep 1 to 26 shadow window	< 1 min	pilot exit (Oct
Response promises kept GUARDRAIL · BUSINESS	Share of jobs meeting the response window promised in the customer's contract	93.5%	≥ 93.5%	every month

DECISION: WHAT EXACTLY TRIPS A GUARDRAIL

Assignment speed worsening more than **15% against its baseline over seven consecutive days**, or response promises falling **below 93.5% in any calendar month**. A trip pauses rollout to further hubs; the dispatch steering group of J. Whitfield (operations sponsor), A. Reyes (product lead), and T. Fontaine (account management) decides within one business day to continue, adjust, or roll back.

EVIDENCE: BASELINES ARE CAPTURED, NOT PROMISED

Idle rate: telematics weekly exports, May to June 2026. Assignment speed: measured on the legacy process during the **September 1 to 26 shadow window**, same timestamp

method as the June sample, owner P. Novak (operations analyst); the baseline exists before the board does. Response promises: existing contract-compliance report.

The events feeding these numbers (a job created, a recommendation shown, a dispatcher choosing differently, an assignment accepted, a job completed) are wired in **before development starts**, and the dashboard exists **before launch day** with the baseline already on it. Precise definitions live in the [Measurement Annex](#).

§5 Scope Boundary, Timeline & Non-Goals

mandatory

2 decisions

1 evidence rule

What this project delivers, when each phase lands, and, just as binding, what it refuses to deliver. A clear boundary is what lets the team say no to good ideas that would sink the date.

In scope

One live dispatch board for all five Eastern-region hubs; skill-and-proximity assignment recommendations; a technician day view on mobile; an exception dashboard for operations managers (second slice).

Delivery frame

Aug to Sep 2026 Build · shadow trial of the distance proxy (Sep 1 to 14)	Sep 1 to 26 Assignment-speed baseline captured on the legacy process	Oct 1 Toronto pilot cutover	Nov 3 to Dec 12 Remaining four Eastern hubs, gated by guardrails
Jan 2027 Second slice: exception dashboard, fairness reporting		Mar 31, 2027 North Star target evaluated	

Non-goals

DEFERRED

Route optimization across multi-day schedules

Belongs to the Routing initiative; owner R. Okafor (VP Operations), scoping begins Q1 2027. Keeping it out protects the October 1 date; it also explains the 1-point gap to the industry benchmark (§2).

REJECTED

Customer-facing technician tracking page

Declined by R. Okafor on June 20, 2026; it requires a privacy review that is a project of its own ([decision record^{\[5\]}](#)).

DEFERRED

Billing and invoicing integration

Waits for the finance system replacement (scheduled Q3 2027); idle-time measurement does not depend on it.

§6 Solution Overview & Key Workflows

mandatory

2 decisions

1 evidence rule

flows render as diagrams


The shape of the solution: detailed enough to understand the approach, deliberately not detailed enough to dictate the design. The primary workflow is a **BPMN 2.0** diagram (the industry-standard process notation) produced through the same rendering service Specira's export pipeline uses; the text below each diagram is its accessibility fallback.

Everything happens on one live board. Jobs arrive on the left; technicians appear on a map and a timeline on the right. When a new job lands, the board proposes the three best candidates (qualified for the work, close to the site, with room in their day), each with a plain-language reason ("certified for commercial HVAC, 12 minutes away, free until 3pm"). Dana accepts a suggestion or picks someone else; Marcus gets the job on his phone; Priya's exception view updates by itself.

DECISION: APPROACH CHOSEN, ALTERNATIVES DECLINED

Recommendation inside the dispatcher's flow was chosen over two alternatives: full auto-assignment (declined: trust, and the union constraint tracked as the open question in §13) and a nightly batch optimizer (declined: doesn't help same-day work, which is 70% of volume per the June job-mix sample^[11]).

Primary workflow: new job to accepted assignment



BPMN 2.0 (lanes: customer, dispatcher, board, technician); rendered by bpmn-js through Specira's diagram service, the same path agents use for exports. Fallback: a job arrives →

Dana records it → if suggestions arrive on time the board proposes three qualified, nearby technicians with reasons (late: the full list shows instead) → Dana accepts or overrides with a reason → the assignment is recorded → Marcus confirms from his phone.

Supporting workflows

- No-show reassignment**

Marcus can't reach a site → he flags it from his phone → the job returns to the board flagged urgent → Dana reassigns with the same suggestions, now excluding him.
- End-of-day exception review**

Priya opens the exception dashboard → sees jobs at risk of missing their promised response window, with causes → follows any job into its full history. (Second slice, January 2027.)

RABBIT HOLE: CALLED OUT AND SIMPLIFIED

Real-time traffic routing is a project-killer if attempted. This project uses straight-line distance plus historical drive times instead, good enough for ranking three candidates. Validated by the September 1 to 14 shadow trial before build commits (§11).

§7 Feature & Requirement Inventory

mandatory 2 decisions 1 evidence rule

Every feature, ranked, with a one-line reason for its rank, and a link to the document that specifies it fully. This PRD never restates what those documents say; it tells you where they are.

Feature	Priority	Slice	Why this rank	Specified in
Live dispatch board	Must have	Pilot (Oct 1)	Every workflow runs through it	Business requirements: Dispatch board^[1]
Assignment recommendations	Must have	Pilot (Oct 1)	Carries the idle-time business case	Functional spec: Recommendation rules^[2]
Technician day view (mobile)	Must have	Pilot (Oct 1)	Without it, assignments still end in phone calls	Business requirements: Dispatch board^[1]

Feature	Priority	Slice	Why this rank	Specified in
Exception dashboard	Should have	Second slice (Jan 2027)	High value for Priya; the board alone already improves idle time	Story-level spec due Sep 15, 2026 · A. Reyes
Fairness reporting	Should have	Second slice (Jan 2027)	Makes the §2 technician protection inspectable before rollout completes	Story-level spec due Sep 15, 2026 · A. Reyes

DECISION: PRIORITIZATION

Ranked by J. Whitfield (operations sponsor) and A. Reyes (product lead) on July 2, 2026, against one question: does removing this feature break the idle-time case? "Must have" means yes. Fairness reporting was raised from Could-have to Should-have on July 15; it is the inspectable form of the §2 protection commitment.

§8 Experience & Quality Constraints

mandatory 3 decisions 1 evidence rule

How the product must feel to the people using it. Each line has a precise engineering threshold behind it; those numbers live in the functional specs in this package, and each constraint links to its home. Every constraint has one named owner.

The product must...	You'd notice a violation when...	Owner	Threshold lives in
Feel live: a new job appears on the board before the dispatcher finishes the intake call	Dana refreshes the page to see if something new arrived	E. Sandoval (engineering lead)	Functional spec: Board refresh, in this package^[3]
Stay usable during a bad connectivity day at a hub	Dispatchers revert to the phone because the board feels unreliable	E. Sandoval	Functional spec: Board refresh^[3]
Be fully usable with a screen reader (WCAG 2.1 AA)	The Toronto or Halifax center, both of which use assistive technology per	D. Kaur (design lead)	Functional specs, accessibility section

The product must...	You'd notice a violation when...	Owner	Threshold lives in
	the discovery records ([8]), can't complete an assignment unaided		
Never lose an accepted assignment	A technician and a dispatcher disagree about who was assigned what	E. Sandoval	Functional spec: Recommendation rules^[2]

§9 UX Reference & Screen Inventory

conditional · active: this project has user-facing surfaces

1 decision

1 evidence rule

The surfaces users will touch, and where their designs live. This section appears because Dispatch Modernization ships user interfaces; a back-office integration project would omit it, with the reason recorded.

Surface	Status	Primary user	Design reference
Dispatch board	New	Dana	Wireframes: Dispatch board, v3^[4]
Technician day view (mobile)	New	Marcus	Wireframes: Dispatch board, v3^[4]
Exception dashboard	New	Priya	Second slice: wireframes due Sep 15, 2026, D. Kaur
Job intake form	Modified	Dana	Wireframes: Dispatch board, v3^[4]

All surfaces follow Meridian's existing design system; this project introduces no new visual language. Each primary workflow in §6 maps to at least one surface above; that mapping is checked by the template's validators.

§10 AI Behavior & Evaluation Contract

conditional · active: this project ships AI-assisted features

3 decisions

1 evidence rule

What the intelligent part of the product is allowed to do, how we'll know it's doing it well, and what happens when it's wrong. Written so a sponsor can hold the team to it, no data-science

background required.

How much the AI decides. At launch, the recommendation is advisory: Dana accepts every assignment herself. Nothing is assigned automatically until the union question in [§13](#) is resolved, and any later increase in automation is its own decision, not a quiet default.

How we'll judge it

Quality we care about	In plain terms	Bar at launch (Oct 1)	Target	By when
Usefulness of suggestions	How often Dana accepts one of the three proposed technicians	≥ 80%	≥ 90%	Mar 31, 2027
Explainability	Every suggestion carries a reason a dispatcher recognizes as true	always	always	continuous
Responsiveness	Suggestions appear within two seconds for nineteen of twenty requests; the engineering threshold lives in the functional spec^[2]	2 s	1 s	Mar 31, 2027

When Dana picks someone other than the suggestion, the board asks her why: one tap, four choices. A. Reyes (product lead) reviews the override reasons in the **monthly operations review**, and they feed the improvement plan; the system learns from its misses on a stated cadence, owned by a named person.

When it fails

If suggestions are **wrong**, Dana overrides; her choice always wins, and the override is recorded. If they are **slow**, the board shows the full technician list immediately rather than making her wait. If the recommendation service is **unavailable**, the board falls back to the last-received technician list it keeps locally and manual assignment against it, visibly flagged as degraded; the mechanism, not just the promise, is specified in the [functional spec^{\[2\]}](#).

DECISION: CAPABILITY, NOT VENDOR

Requirements here are stated as capabilities (ranking with explanations, sub-call-time response). Which AI supplier or model provides them is a configuration choice owned by

engineering, revisitable without reopening this document.

§11 Facts, Assumptions & Risks

mandatory2 decisions1 evidence rule

Three different things, kept honestly apart: facts we've verified, assumptions we're standing on (each with a dated plan to test it), and risks we've decided to carry (each with a named owner who accepted it).

FACT Vehicle tracking covers 100% of the fleet
Verified against the fleet inventory on June 30, 2026; idle-time measurement needs no new hardware.

ASSUMPTION Straight-line distance plus historical drive time is a good-enough proxy for real travel time
Load-bearing: the recommendation ranking depends on it. Validated by the **September 1 to 14 shadow trial**, owner E. Sandoval ([trial requirement^{\[9\]}](#)); build does not commit to pilot until the trial report is accepted. Fallback if it fails: ranking by skills and calendar alone.

ASSUMPTION Dispatchers will record override reasons if it costs one tap
Tested in Toronto pilot weeks 1 to 2, owner M. Chen; if reasons aren't recorded, the monthly learning loop loses its input and the reason-capture design is revisited before hub rollout.

RISK · ACCEPTED Toronto or Halifax dispatch teams may resist suggestion ordering
Mitigation: suggestions are advisory and overriding is one click with no friction. Accepted by R. Okafor (VP Operations), July 2, 2026.

§12 Dependencies & Mandated Constraints

mandatory1 decision1 evidence rule

The rules this project didn't choose but must obey, each with the authority that mandates it, and the outside parties it depends on. A "constraint" without a citable authority gets reclassified as an assumption; that rule is enforced by the template.

Constraint or dependency	Kind	Authority / owner	If it bites
Union agreement governs how work	Mandated	Collective agreement, §14	Automation stays at "recommendation only"; see §13

Constraint or dependency	Kind	Authority / owner	If it bites
may be ordered and assigned			
All customer and location data stays in Canada	Mandated	Enterprise customer contracts	Constrains hosting choices; already satisfied by current infrastructure
Vehicle-tracking data feed	Dependency	S. Grewal (fleet manager; external vendor relationship)	Idle-time measurement falls back to manual timesheets with a one-week delay; the North Star goes blurry, the product keeps working

§13 Open Questions

mandatory

1 decision

What we don't know yet: each question with a person responsible for answering it and a date. An unowned question is not tracked; that rule is enforced by the template.

Question	Owner	Answer by	What it blocks
Does the union agreement constrain automated assignment ordering?	N. Duval (Legal & Operations lead)	Aug 1, 2026	How much the AI may automate: <u>§10</u>
Will the tracking-feed vendor commit to a same-day data guarantee?	S. Grewal (fleet manager)	Aug 15, 2026	Nothing at launch; affects how fresh the idle-time number can be

RISK: IF THE UNION QUESTION IS UNRESOLVED BY ITS DATE

Assignment ships recommendation-only (Dana approves every assignment) until the constraint is cleared. Accepted by R. Okafor (VP Operations) on July 2, 2026. The **October 1 pilot date does not move**; only the automation posture does.

Refs References & Package Contents

Everything this document points at, in one place. Package items travel in the same export bundle as this PRD; workspace items link into Specira, where the full record and its history live.

In this export package

PACKAGE	[1] Business requirements: Dispatch board ./brd-dispatch-board.docx
PACKAGE	[2] Functional spec: Recommendation rules ./frd-recommendation-rules.docx
PACKAGE	[3] Functional spec: Board refresh & resilience ./frd-board-refresh.docx
PACKAGE	[4] Wireframes: Dispatch board, v3 ./wireframes-dispatch-board-v3.html

In the Specira workspace

SPECIRA	[5] Decision: Customer tracking page declined (R. Okafor, June 20, 2026) app.specira.ai/projects/dispatch-modernization/decisions/86
SPECIRA	[6] Discovery interview: Dispatcher, Toronto hub (June 12, 2026) app.specira.ai/projects/dispatch-modernization/discovery/sessions/14
SPECIRA	[7] Knowledge base: Exit-interview themes, Q2 2026 app.specira.ai/kb/entries/exit-interview-themes-2026q2
SPECIRA	[8] Personas & discovery records: Toronto and Halifax hubs (May to June 2026) app.specira.ai/projects/dispatch-modernization/personas
SPECIRA	[9] Requirement: Distance-proxy shadow trial, Sep 1 to 14, 2026 app.specira.ai/projects/dispatch-modernization/requirements/RQ-041
SPECIRA	[10] Fact: Idle-cost analysis, \$2.1M/yr total, \$700k/yr assignment-delay share app.specira.ai/projects/dispatch-modernization/facts/idle-cost-analysis
SPECIRA	[11] Fact: Job-mix sample, 70% same-day (June 2026) app.specira.ai/projects/dispatch-modernization/facts/job-mix-jun26

Annex Measurement & Instrumentation Annex

The technical identifiers behind the plain-language names used in this document, for the engineers and analysts who wire them up. Readers of the body never need this table; auditors and implementers do.

Plain name in this document	System identifier	Source & cadence
Technician idle rate	technician_idle_rate	telematics feed · weekly · 4-week trailing for the target test
Assignment speed	dispatcher_assignment_time_p75	product events · daily · baseline from Sep 1 to 26 shadow window
Response promises kept	sla_breach_rate (inverted)	job records · monthly
Instrumented events	job_created · recommendation_shown · recommendation_overridden · assignment_accepted · job_completed	named before build

HOW MISSING CONTENT RENDERS: CONTRACT STUBS AND OMISSIONS

The example above is complete. When a section has no content yet, reviewers see the template's demands instead of a blank heading, and when a conditional section legitimately doesn't apply, the document says so instead of silently skipping it:

\$10 AI Behavior & Evaluation Contract

conditional · AI features

no content yet

This section still needs: evaluation criteria in three classes with launch and target bars · risk tiers with safeguards per tier · what happens when the AI is wrong, slow, or unavailable · capability needs stated without vendor names. 3 decisions and 1 evidence rule outstanding.

Omission note (how a skipped conditional section renders): "\$9 UX Reference omitted: back-office integration project with no user-facing surfaces. Rationale recorded in adaptation event #3."

RENDERER SPEC: PRESENTATION PROFILE TOKENS

primary

ink

muted

#0D9488

#0F172A

#475569

surface
#F8FAFC

accent (risk only)
#F59E0B

Type: Inter (headings + body, base 11pt in paged output), JetBrains Mono for hashes, data columns, and reference targets, always tabular numerals. Callouts: decision = teal left border · evidence = slate · risk = amber · omission = muted italic in a dashed frame. Tables: teal header, 4% teal zebra. Cross-references: teal dotted underline with a superscript reference number resolving through the References list. Diagrams: flows render as images produced by Specira's diagram service (BPMN-style swimlanes for business process, sequence for interaction, state for lifecycle) on white cards, each with a textual accessibility fallback beneath it. Every section opens with an italic reader-guidance line. System identifiers never appear in body prose; they live in the Measurement Annex. Contract chips under section heads expose the governed demands in review renderings; final stakeholder exports may suppress them via output profile configuration.