GOVERNED TEMPLATE RENDERING REFERENCE: THREAT MODEL DEFAULT V2 (DRAFT)

definition e8ac9299…a660

THREAT MODEL · PROJECT ARTIFACT

SPECIRA

Threat Model: Dispatch Modernization

Meridian Field Services: instructional example, not project evidence

DRAFT · WATERMARK POLICY: DRAFT_ONLY

template threat_model v2 · pack: specira_default_delivery

analyze here, remediate by pointer; SR rows own the fixes

§1 Scope & Assumptions

mandatory | 1 decision | 1 evidence rule

validators: every_threat_cites_architecture_element_or_crossing · no_architecture_diagram_reproduced

Q1 of the four-question spine: what are we building. The architecture's views are walked by reference, never redrawn; every assumption is a challengeable claim, not scenery.

System, by reference: the container view, the data-flow view with its marked crossings, and the integration register (architecture[1]). **Trust boundaries, sourced**: `PB-1` feed ingress (personal location data crosses in), `PB-2` provider boundary (anonymized codes only), the public mobile surface, and the transitional `INT-3` file drop.

Assumption (each challengeable)	If it falls
The cloud platform's tenant isolation holds	Re-model with the platform inside the boundary
Meridian SSO is out of scope, owned by the identity team (auth & authz policy[2])	Federation threats join the register
Attacker ceiling: motivated external actor or malicious insider with dispatcher rights; not a state actor	Ratings scale re-anchors
The vendor's signing keys are not already compromised	`TM-01` controls are void; incident path

Scope this iteration: the pilot system as designed; cutover changes re-model on their named trigger (§7).

§2 STRIDE Analysis

mandatory · 1 decision · 1 evidence rule

validators: every_crossing_analyzed_or_explicitly_waived · attacker_story_present_and_non_generic · disposition_mitigate_cites_sr_id_or_architecture_decision

Per-interaction sweep over every marked crossing; per-element depth on the event log and position cache (regulated classes). Each row is an attacker story with a rating from §5's anchored scales and a disposition that points; remediation text lives in the security requirements, restated nowhere.

Id	Class	Element / crossing	Attacker story	Existing controls	Rating
TM-01	Spoofing	INT-1 feed at PB-1	An actor replays or forges vendor webhooks to poison positions and steer assignments	Signature verification exists	likely · high personal class crosse
TM-02	Tampering	Job intake	Malformed or oversized skill payloads corrupt the queue or smuggle content into board rendering	none	possible · moderate
TM-03	Elevation	Board role boundary	A read-only operations account crafts assignment calls directly against the	none	possible · high

Id	Class	Element / crossing	Attacker story	Existing controls	Rating
			interface's actions		
TM-04	Info disclosure	PB-2 provider boundary	Prompt-shaped content or over-broad inputs leak location or identity to the model provider	Anonymized codes only, by design	possible · high personal class
TM-05	Repudiation	Override actions	A dispatcher disputes an assignment they made	Append-only audit	rare · moderate
TM-06	Denial	INT-2 notification path	Flooding or gateway failure silently drops assignment delivery	Degradation posture designed	possible · moderate

Coverage: every crossing from §1 resolves to rows. INT-3 is **waived** for STRIDE depth beyond tampering, with rationale: transitional, retired at cutover, integrity-checked by re-run (INT-3[1]); **waiver expires at the cutover re-model**.

§3 Privacy & AI Threats

conditional · both triggers true, ENGAGED | 1 decision | 1 evidence rule

validators: privacy_threats_present_when_personal_data_trigger · ai_threats_present_when_ai_signal_true

STRIDE misses what privacy and AI taxonomies catch: personal data flows engage LINDDUN classes; the AI boundary engages over-trust and exfiltration.

Id	Taxonomy · class	Threat	Disposition
TM-07	LINDDUN · identifying, linking	Position tracks re-identify technicians even without names: detectability of shift patterns	mitigate → retention and residency rows (SR-033, SR-034[3]) + the 30-day purge (data architecture[1])
n/a	LINDDUN · unawareness	Dispatcher override logging touches worker monitoring; the consent-notice question is open	tracked → security open questions[3]
TM-08	AI · over-trust	The board acting on a hallucinated candidate	**eliminate**: structurally, by the constrained-output requirement (SR-016[3] citing FR-013[6]), decision ADR-5[7]
n/a	AI · exfiltration	Carried by `TM-04` above: one threat, one row, linked not duplicated	see `TM-04`

§4 Attack Surface

mandatory · 1 decision · 1 evidence rule · validator: attack_surface_orders_by_blast_radius

Ordered by blast radius, not alphabetically; the first row is where an hour of attacker time buys the most.

Entry point	Reachable from	Auth gate	Data classes reachable	Threats
Public mobile surface (technician clients)	Internet	SSO-gated	Own assignments	TM-05, TM-06
Vendor webhook endpoint	Partner-reachable	Signed tokens	Positions	TM-01
Dispatch board	Staff network	SSO + roles	Jobs, assignments, positions	TM-02, TM-03
Nightly INT-3 drop (transitional)	Internal	none	Job master	waived, with expiry (§2)

Reconciliation: no surprise entry points against the architecture's context view (context[1]); checked, clean.

§5 Risk Ratings

mandatory | 1 decision | 1 evidence rule

validators: likelihood_impact_from_anchored_scale · impact_anchors_cite_data_classes · acceptance_floor_has_owner

Anchored scales set once, with the owners named; a rating without an anchor is a mood. Set with N. Duval and E. Sandoval.

Likelihood	Anchor
rare	Under 10%: no precedent in comparable fleets
possible	10 to 50%: a plausible path exists
likely	Over 50%: recurred in comparable systems

Impact	Anchor (cites data classes)
low	No regulated class; single-user annoyance
moderate	Confidential class exposed, or one hub degraded
high	Personal or regulated class exposed, or a marked boundary crossed uncontrolled; classes per data protection[3] until the classification register ships (data classification[8])

Why not the alternatives: per-factor scoring retired for rater subjectivity; vulnerability-severity scores answer a different question (a bug's severity, not a design threat's risk). **Acceptance floor**: rare-low and rare-moderate may be accepted by the engineering lead; anything touching a personal class escalates to N. Duval. **Material ratings** feed the risk register by id (risk register[9]).

§6 Mitigation Strategies

mandatory | 1 decision | 1 evidence rule

validators: disposition_from_closed_enum · no_security_requirement_text_restated · handshake_resolves_both_directions

Dispositions roll up; the text of every fix lives in exactly one place: the SR row or the architecture decision the pointer names.

Disposition	Count	Rows	Where the fix lives
mitigate	6	TM-01, TM-02, TM-03, TM-04, TM-06, TM-07	Each points at its SR rows or architecture

Disposition	Count	Rows	Where the fix lives
			decision; restated nowhere
eliminate	1	TM-08	The constraining decision ADR-5[7] removes the path structurally
accept	1	TM-05 residual	A dispatcher disputing a logged action they performed; accepted rare-moderate by **N. Duval**; revisit at the union answer (open questions[3])
transfer	0	none	None this iteration

Handshake, checked both directions: every TM above the acceptance floor resolves to SR rows; every SR driver citing a TM-id resolves back. Verified at generation; no dangling ids.

§7 Retrospective & Revisit Triggers

mandatory · 1 decision · 1 evidence rule

validators: revisit_triggers_named_not_generic · walk_history_carries_dates_and_architecture_ref

Did we do enough? Answered honestly, plus the named events that reopen the model. A one-time workshop PDF is the anti-pattern this section exists to kill.

Coverage judgment: crossings and the AI boundary covered well; the mobile client's device-level threats covered thinly, and deliberately: the field app extension inherits the field-tools team's model, linked (sourcing · extend[1]); `INT-3` deferred behind its waiver expiry.

Revisit trigger (named)	What reopens
The cutover re-model: INT-3 retirement + job authority transfer	Full walk; the waiver expires
Any new integration joining the register	Its crossing's rows
A new data class entering the system	Impact anchors + affected rows

Revisit trigger (named)	What reopens
The union answer changing audit scope	TM-05's acceptance
Any incident touching a boundary	That boundary's rows + the assumption it broke

Walk history: modeled July 14, 2026 against the architecture stamped 2026-07-16, by E. Sandoval with N. Duval; next scheduled walk at cutover.

§8 Open Questions

mandatory 1 decision

Question	Owner	Answer by	Blocks
Does worker-monitoring guidance require a dispatcher-facing notice for override logging (the LINDDUN unawareness row)? Shared with the security requirements' question, tracked once, linked twice (open questions[3])	N. Duval	Aug 8, 2026	TM-05's acceptance revisit only

Refs References & Package Contents

In the Specira workspace

specira [1] Architecture: container/data-flow/context views, integration register, INT-2/INT-3 semantics, sourcing app.specira.ai/projects/dispatch-modernization/artifacts/architecture

specira [2] Auth & authz policy: federation scope, identity ownership app.specira.ai/projects/dispatch-modernization/artifacts/auth-authz-policy

specira [3] Security requirements: SR rows cited by disposition pointers; open questions app.specira.ai/projects/dispatch-modernization/artifacts/security-requirements

specira [4] User stories: US-1 acceptance row AT-5 (server-side denial) app.specira.ai/projects/dispatch-modernization/artifacts/user-stories#us-1

specira [5] Deployment verification: DV-5 notification-path check app.specira.ai/projects/dispatch-modernization/artifacts/deployment-verification#dv-5

specira [6] FRD: FR-013 constrained candidate output app.specira.ai/projects/dispatch-modernization/artifacts/frd#fr-013

specira	[7] Decision log: ADR-5 constrained-output decision app.specira.ai/projects/dispatch-modernization/decisions/adr-5
specira	[8] Data classification: label register (impact anchors re-cite on ship) app.specira.ai/projects/dispatch-modernization/artifacts/data-classification
specira	[9] Risk register: material ratings land by TM-id app.specira.ai/projects/dispatch-modernization/artifacts/risk-register

Generated by Specira · template threat_model v2 (draft) · pack specira_default_delivery lineage e8ac9299…a660 · page 1 of 6