

GOVERNED TEMPLATE RENDERING REFERENCE: OBSERVABILITY & MONITORING DEFAULT V2 (DRAFT)

definition aa5edf72...3897

OBSERVABILITY & MONITORING · PROJECT ARTIFACT

SPECIRA

Observability & Monitoring: Dispatch Modernization

Meridian Field Services: instructional example, not project evidence

DRAFT · WATERMARK POLICY: DRAFT_ONLY

template observability_monitoring v2 · pack: specira_default_delivery

measure commitments, never mint targets; the architecture names, this designs

§1 Logging Strategy

mandatory 1 decision 1 evidence rule

validators: log_pii_fields_cite_classification_label · pii_redaction_at_source_stated · personal_retention_shorter_than_audit_retention

Structured logs joined to metrics and traces by a propagated correlation id. PII is a classification concern: redacted at source, nothing downstream to scrub.

Standard: structured JSON; required fields are timestamp, level, service, correlation id; the id is issued at ingress and propagated through the queue to the worker: one id joins a position's webhook arrival to its board render.

Field family	Label (cites)	Log rule
Position.lat / lon	RESTRICTED (labels ^[1])	Never enter logs: the adapter logs band ids only; redaction at source
Technician.name	CONFIDENTIAL	Logs nowhere; hr_id appears only in the audit stream

Retention tiers: hot 30 days (incident queries); warm 12 months (trends); cold archive. The auth audit stream retains 2 years per the retention rules^[1], longer than the operational logs' personal-data window, as the register requires.

§2 Metrics, SLOs & Dashboards

mandatory

1 decision

1 evidence rule

validators: every_slo_cites_nfr_or_quality_scenario ·
every_architecture_named_signal_instrumented · every_dashboard_answers_named_question

The instrumentation standard is non-optional; SLOs measure what was committed; a dashboard without a question is wallpaper and does not ship.

Standard, every service: request rate, error count, duration histograms, the correlation id; the adapter and worker add queue depth and batch latency. **Architecture handshake:** every signal the architecture named observable is instrumented by citation ([observability signals^{\[2\]}](#)): queue depth, adapter batch latency, dead-letter rate, database write saturation.

SLI	SLO (cited)	Error budget
Board liveness: job visible	≤ 5 s p95 (NFR-P1^[3])	0.5% of dispatch-hours minutes
Suggestion latency	≤ 1 s p95 (NFR-P2^[3])	same window
Availability	99.5% over dispatch hours (NFR-A1^[3])	≈ 24 min / month

No target minted here.

Dashboard	The question it answers	Audience
Dispatch health	"Is dispatch healthy?": the four golden signals per service	On-call
Where next	"Where do I look next?": queue depth, feed gap, dead-letter trend	On-call
Value case	"Is the value case holding?": override rate, suggestion acceptance (RSK-06^[4])	M. Chen, weekly

No fourth dashboard; wallpaper does not ship.

§3 Alerting & On-Call

mandatory

1 decision

1 evidence rule

validators: every_alert_has_owner_and_runbook_link · every_alert_has_first_action ·
page_vs_ticket_severity_explicit_per_alert · kris_cited_by_risk_id

Every alert actionable, owned, runbook-linked, with a first action. An alert that cannot state its first action is deleted, not shipped.

Alert · threshold	Severity	Owner	First action	Runbook / KRI
Feed gap ≥ 15 min	PAGE	On-call engineer	Confirm vendor status, then activate manual assignment (INT-1 response ^[5])	feed-loss runbook ^[6]
Signature failures > 1% / 10 min	PAGE	On-call engineer	Inspect replay-window rejections; rotate on confirmation (SR-021 ^[7])	KRI: RSK-01 ^[4]
Dead-letter rate > 0 for 15 min	ticket	Platform engineer	Inspect the poison message class	none
Override rate > 40% after week 2	ticket · weekly review	M. Chen	Open the override-reason review loop	KRI: RSK-06 ^[4]
Suggestion degradation sustained 30 min	ticket	Platform engineer	Check provider status against the fallback contract (suggestions_unavailable ^[8])	none

Fatigue: monthly review deletes unowned or never-actioned alerts; precision tracked as actionable-over-total, floor 70%. **On-call:** one primary, one secondary, weekly rotation; escalation to E. Sandoval defined before the first page; toil budget under 30% of on-call time; two consecutive weeks above it escalates to management, not heroics.

§4 Open Questions

mandatory

1 decision

Question	Owner	Answer by	Blocks
Does the value-case dashboard need dispatcher-visible variants (the union briefing may require transparency about what is measured; open questions ^[7])?	M. Chen with N. Duval	Aug 22, 2026	The dashboard register's audience column only

Refs References & Package Contents

In the Specira workspace

specira	[1] Data classification: labels, retention rules app.specira.ai/projects/dispatch-modernization/artifacts/data-classification
specira	[2] Architecture: the named observability signals app.specira.ai/projects/dispatch-modernization/artifacts/architecture#observability
specira	[3] NFR catalog: NFR-P1, NFR-P2, NFR-A1 (the committed targets) app.specira.ai/projects/dispatch-modernization/artifacts/nfr-catalog
specira	[4] Risk register: RSK-01, RSK-06 (the KRIs wired here) app.specira.ai/projects/dispatch-modernization/artifacts/risk-register
specira	[5] Integration inventory: INT-1 operational response app.specira.ai/projects/dispatch-modernization/artifacts/integration-inventory#int-1
specira	[6] Deployment verification: feed-loss runbook app.specira.ai/projects/dispatch-modernization/artifacts/deployment-verification
specira	[7] Security requirements: SR-021; shared open question app.specira.ai/projects/dispatch-modernization/artifacts/security-requirements
specira	[8] API contracts: the degradation contract app.specira.ai/projects/dispatch-modernization/artifacts/api-contracts#errors