GOVERNED TEMPLATE RENDERING REFERENCE: PERFORMANCE TEST SCENARIOS DEFAULT V2 (DRAFT)

definition 19d6ba65…df8f

PERFORMANCE TEST SCENARIOS · PROJECT ARTIFACT

SPECIRA

Performance Test Scenarios: Dispatch Modernization

Meridian Field Services: instructional example, not project evidence

DRAFT · WATERMARK POLICY: DRAFT_ONLY

template performance_test_scenarios v2 · pack: specira_default_delivery

instantiate, never invent; targets cite the architecture

§1 Scope & Committed Targets

mandatory | 1 decision | 1 evidence rule

validators: criteria_cite_architecture_scenario_ids_no_bare_numbers · whole_artifact_omits_when_no_committed_targets

Every number below was committed in the architecture and is cited by id. Were no targets committed, this artifact would not exist; load-testing against invented numbers is theater.

Instantiates: board liveness[1] (job visible ≤5 s p95), feed-loss resilience[1] (zero blocked assignments), and the scale frame: five hubs, 1,500 jobs/day, 400 positions/minute, **first bottleneck expected at adapter fan-in beyond 2,000 positions/minute**: restated by reference, never re-derived. Scope this phase: liveness, fan-in, burst; agreed with E. Sandoval.

§2 Load Scenarios

mandatory | 1 decision | 1 evidence rule

validators: workload_model_required_before_rates · scenario_has_resolvable_script_ref · criteria_generated_or_stamp_verified_against_script

Model before rates, with the reason it matches real traffic. The script is the executable truth; criteria here carry the verified-at-commit stamp against its threshold block.

Scenario	Type	Model (reason)	Rate / ramp / duration
PS–1 board liveness	baseline + load	**open** arrival: jobs arrive from the world; dispatchers do not queue behind each other	4 jobs/min · ramp 5 min · steady 1
PS–2 shift-start spike	spike	**open**: the 7 a.m. burst is arrival-shaped (observed pattern[6])	burst 20 jobs/min · 10 min over ba
PS–3 adapter fan-in	breakpoint	**open** arrival ramp: vendor push shape	400 → 2,500 positions/min · 300 s
PS–4 pilot-day soak	soak	**open** at average load	4 hours steady

§3 Test-Type Coverage

mandatory · 1 decision · 1 evidence rule

validator: conditional_types_link_risk_flag_or_omit_rationale

Type	Disposition	Basis
Baseline + load	**always**: PS-1 carries both at this scale	n/a
Spike	engaged · PS-2	The shift-start burst is the observed pattern
Breakpoint	engaged · PS-3	The architecture names a first-bottleneck hypothesis to settle; breakpoint finds the ceiling, a different question from stress
Soak	engaged · PS-4	Position cache + pooled connections are a named leak risk (caching decisions[1])
Stress	omitted, with rationale	Graceful degradation past normal load is already exercised by the feed-loss failure-injection suite (risk prioritization[3]); redundancy is the rationale; "slow" would not be

§4 Test Environment & Volumetrics

mandatory · 1 decision · 1 evidence rule

validators: environment_parity_and_sizing_stated · parity_gaps_named_with_owners

Parity: pilot topology (production-shaped, Toronto scale), same container layout; managed PostgreSQL **one tier below production sizing, a named gap**, accepted by E. Sandoval; it limits the breakpoint claim to a relative, not absolute, ceiling; re-run at production sizing before hub-six rollout (open questions[1]).

Data volumetrics: 500 thousand position rows preloaded (the 30-day working set), 1,200 open-and-recent jobs, 300 technicians, matching the data model's volumetric note[2]; an empty database predicts nothing. **Ladder**: the strategy's pilot tier (environments[3]), cited.

§5 Bottleneck Hypotheses

mandatory · 1 decision · 1 evidence rule

validator: bottleneck_matches_architecture_claim_verbatim

Hypothesis (inherited verbatim): first bottleneck at adapter fan-in beyond 2,000 positions/minute (scale frame[1]). **Instrumentation**: queue depth, adapter batch latency, dead-letter rate, database write saturation (observability signals[1]). **If wrong**: a database-first ceiling re-scopes PS-3's step profile and fires the architecture's revisit trigger; a higher ceiling raises confidence in the hub-six answer.

§6 Results, Interpretation & Re-run Triggers

mandatory · 1 decision · 1 evidence rule

validators: results_use_percentiles_never_means · result_tables_cite_run_and_commit

Percentiles, never means; the tail is the user experience. Each run settles or reshapes the hypothesis; distributions attach to the package.

Scenario · run	Target	Actual (p95 / p99)
PS-1 · 2026-09-12 (commit-stamped)	≤ 5 s p95	3.8 s / 4.6 s
PS-3 · 2026-09-12	hypothesis: 2,000/min	failure condition at 2

Full distributions and the soak time series attach (perf run exports[4]), referenced never inlined. **Re-runs**: before the release gate (CI gates[3]), after any scale-frame change, monthly during hub rollout.

§7 Open Questions

mandatory · 1 decision

Question	Owner	Answer by	Blocks
Does the breakpoint re-run at production database sizing move before or after the hub-six decision?	E. Sandoval	Aug 15, 2026	The absolute-ceiling claim only

Omission note (no committed targets): "performance_test_scenarios omitted for this project: the architecture commits no numeric performance targets; load-testing against invented numbers is theater. Recorded in adaptation event #1."

Refs References & Package Contents

In this export package

package	[4] Perf run exports: distributions, soak time series ./perf/run-exports/

In the Specira workspace

specira	[1] Architecture: quality scenarios, scale frame, caching risk, observability app.specira.ai/projects/dispatch-modernization/artifacts/architecture
specira	[2] Data model: volumetric notes app.specira.ai/projects/dispatch-modernization/artifacts/data-model#indexing
specira	[3] Test strategy: environments ladder, gates, risk prioritization app.specira.ai/projects/dispatch-modernization/artifacts/test-strategy
specira	[5] CI: perf-nightly job app.specira.ai/projects/dispatch-modernization/ci/perf-nightly
specira	[6] Journey maps: the shift-start burst observation app.specira.ai/projects/dispatch-modernization/artifacts/journey-maps

Generated by Specira · template performance_test_scenarios v2 (draft) · pack specira_default_delivery lineage 19d6ba65…df8f · page 1 of 5