Jacquard Router: Tuning and Analysis (Draft)

This report studies Jacquard routing behavior across seven engines using a common simulator corpus and analysis pipeline. The engines in scope are `batman-classic`, `batman-bellman`, `babel`, `olsrv2`, `scatter`, `pathway`, and `field`. The goal is to understand where each engine works well, where performance degrades, what kind of failures arise, and to compare engines under the same network conditions. Routing quality is regime-dependent: a setting that works well in an easy connected network may break down under asymmetry, bridge loss, candidate pressure, or uncertainty.

The report is organized in four parts. Part I covers tuning: recommended configurations, transition behavior, failure boundaries, and simulator assumptions. Part II covers engine-specific analysis and cross-engine comparisons. Part III calibrates diffusion-oriented engine profiles. Part IV evaluates these calibrated profiles under message-diffusion scenarios where node movement and intermittent contact opportunities carry messages, and end-to-end paths may not exist.

Design Setting

The maintained corpus is designed for disrupted and mobility-driven mesh environments. In this setting, end-to-end paths are often absent. Connectivity appears through short contact windows, weak bridges, and repeated partial recovery rather than through one stable connected graph. Nodes are also resource-constrained, so routing quality depends on bounded state, bounded work, and disciplined use of transmissions and custody.

The route-visible matrix gives useful evidence for this setting because it stresses the conditions that determine whether a router-facing engine remains usable at all. The maintained families vary bridge pressure, asymmetry, loss, relink events, partitions, recovery, contention, and local node pressure. Those are the same forces that determine whether a proactive engine keeps a route, whether a search-driven engine finds one, and where each approach breaks down.

The diffusion track adds the second half of the picture. It models cases where movement is the transport mechanism and messages must persist across disconnection. Its mobility-driven contacts, bounded replication, energy and transmission accounting, storage utilization, and observer-leakage measures give insight into whether a deferred-delivery policy remains viable in the same population-level setting, not only in easy connected regimes.

Part I. Tuning

1. Recommended Configurations

This table condenses the highest-ranked configurations for each engine family.

Engine Configuration	Score	Activation	Route Presence	Max Stress
`batman-classic-2-1`	5071.3	1000.0	639.3	48
`batman-classic-4-2`	5070.0	1000.0	653.0	48
`batman-bellman-1-1`	4289.2	1000.0	780.0	56
`batman-bellman-2-1`	4277.7	1000.0	780.0	56
`babel-8-4`	4518.2	1000.0	922.3	52
`babel-6-3`	4499.0	1000.0	922.3	52
`olsrv2-8-4`	5072.7	1000.0	792.2	52

Engine Configuration	Score	Activation	Route Presence	Max Stress
`olsrv2-6-3`	5039.7	1000.0	781.8	52
`pathway-2-hop-lower-bound`	4166.3	1000.0	876.3	58
`pathway-2-zero`	4166.3	1000.0	876.3	58
`scatter-balanced`	5292.4	1000.0	900.5	90
`scatter-conservative`	5271.8	1000.0	900.5	90
`comparison-b6-3-p4-hop-lower-bound`	5253.3	1000.0	895.5	90
`comparison-b4-2-p3-zero`	3333.2	569.4	517.7	42
`field-4-zero-p1-f140-n180`	4606.5	1000.0	780.9	60
`field-8-hop-lower-bound-p1-f120-n190`	4606.5	1000.0	780.9	60

Table 1. Score is the composite ranking value; Activation is the share of runs that installed a route; Route Presence is the average fraction of rounds with a live route; Max Stress is the highest sustained stress level survived before the first maintained breakdown.

Detailed transition, failure-boundary, profile, and `field`-regime tables are collected in Appendix A so the main report can stay focused on the key recommendations and figures.

2. Tuning Setup And Scoring

Simulation Setup

Each experiment run uses the Jacquard simulator to play out a fixed network scenario from a known random seed. A scenario defines the network layout, available routing engines, active routing requests, and round count. The simulator can apply planned changes during a run: cutting links, restoring links, degrading links asymmetrically, moving connections, or imposing local resource limits.

The report scores observable replay output: whether a route appears, when it is first lost, whether it recovers, how often it changes, and what failures are recorded.

The simulator now also records model-lane snapshot and reducer artifacts for validation, but the recommendation tables and figures in this report continue to score the maintained full-stack replay artifacts. The added model-lane output is there to check planner and reducer behavior, not to replace the router-visible outcome surface.

Matrix Design

The tuning matrix changes one small set of conditions at a time. Across the corpus, it varies network density, loss, interference, asymmetry, topology change, node pressure, and objective type. For `batman-classic`, `batman-bellman`, `babel`, and `olsrv2`, the main sweep changes decay-window settings. For `scatter`, the route-visible sweep compares the maintained `balanced`, `conservative`, and `degraded-network` profiles. For `pathway` and `field`, the main sweep changes per-objective search budget and heuristic mode.

The recommendations are meant to be good defaults for this modeled corpus, not single winners from one scenario.

Methodology note:

- route-visible `Active Route` means active-window route presence from objective activation onward, not total scenario-round presence.

- the head-to-head route-visible surface is a fixed representative-profile benchmark, while Part I recommendation tables are calibrated-best surfaces.
- mixed-stack `Selected-Round Leader` means the engine selected for the most active-route rounds inside one shared router policy, not the best standalone engine.
- generic family-by-family diffusion winners are more provisional than the regime-aware diffusion tables when different generic weightings change the winner.

Regime Assumptions

The scenarios are stylized representations of common mesh-network conditions. Names like sparse line, medium ring, and bridge cluster describe the network shape. Loss and interference settings then make communication easier or harder within that shape. Some families are intentionally placed near break points where a small parameter change can determine whether a route survives.

Regime Characterization

Topology regimes:

- `sparse line`: nodes depend on a single chain of relays with few alternate paths.
- `medium ring`: looped connectivity allows one route to fail while another may survive.
- `medium mesh` and `dense mesh`: multiple neighbors can reach the same destination, so contention and search choice matter more.
- `bridge cluster`: two groups joined by one narrow bridge, so a single weak link dominates routing.
- `high fanout`: one node sees many candidate neighbors, stress-testing search budget.

Condition regimes:

- `low`, `moderate`, and `high loss` describe message drop rates.
- `interference` and `contention` describe medium crowding.
- `asymmetry` means one direction of a link is worse than the other.
- `churn`, `relink`, `partition`, and `recovery` describe topology changes over time.
- `intrinsic node pressure` means the node itself becomes a bottleneck.

Workload regimes:

- `connected-only` requires an actual connected route.
- `repairable-connected` allows temporary disruption with expected recovery.
- `service` means the engine chooses among candidate service locations.
- `concurrent mixed` means multiple route requests compete under shared pressure.

BATMAN Classic Algorithm

BATMAN Classic is the spec-faithful BATMAN IV engine: TQ is carried in OGMs and updated multiplicatively at each relay hop, TTL bounds propagation depth, and bidirectionality requires echo-window confirmation. It produces no route candidates until receive-window data has accumulated and echo confirmation has been received. BATMAN Bellman can bootstrap routes from its local Bellman-Ford computation before OGMs arrive, but BATMAN Classic has no such shortcut. The analysis compares its behavior under identical regimes to isolate what the Bellman-Ford enhancement contributes. Its retained state is protocol state only, not deferred payload storage.

BATMAN Bellman Algorithm

BATMAN Bellman tracks a good next hop toward each destination using a local Bellman-Ford computation over a gossip-merged topology graph. It can bootstrap routes from topology data before OGMs arrive. The tuning parameters control how quickly old information expires and how quickly the engine refreshes its view. Like the other conventional proactive next-hop engines in this corpus, it retains routing control state such as learned advertisements and bidirectionality evidence, but it does not buffer payloads for deferred store-and-forward delivery.

Babel Algorithm

Babel implements the RFC 8966 distance-vector protocol. Link cost uses bidirectional ETX rather than forward-only TQ, penalizing asymmetric links more heavily. Path metric is additive rather than multiplicative. Route selection is gated by a feasibility distance table that provides loop freedom during transient topology changes. The analysis targets asymmetric link regimes and partition recovery to surface these behavioral differences. As with the BATMAN variants, the cached state here is advertisement, feasibility, and metric state rather than payloads awaiting store-and-forward delivery.

OLSRv2 Algorithm

`olsrv2` is Jacquard's deterministic proactive link-state baseline. It learns one-hop and two-hop reachability from HELLO exchange, elects a stable MPR covering set, floods TC advertisements only when local topology changes or forwarded MPR-selected TC state is fresher, and runs shortest-path derivation over the learned topology tuples.

The Jacquard engine intentionally simplifies the full RFC surface: it keeps one deterministic link-state view, one MPR election policy, and next-hop-only route publication. The tuning sweep therefore focuses on the decay window that controls how long HELLO and TC evidence stays live and how quickly the engine requests another synchronous round. That retained state is topology and freshness metadata only; `olsrv2` is not a deferred-delivery payload buffer.

Scatter Algorithm

`scatter` is Jacquard's bounded deferred-delivery diffusion engine. It publishes an opaque, partition-tolerant route claim when the objective is supportable, then performs engine-private store-carry-forward movement with hard copy budgets, scope-relative carrier scoring, local regime selection, and contact-feasibility gates.

Unlike an acknowledgement-driven custody protocol, `scatter` does not assume authoritative transfer. It retains bounded payload custody locally, can split copy budget conservatively across better carriers, and may prefer a better carrier without claiming globally reliable handoff semantics. Its retained state is payload custody plus local diffusion metadata, not a topology graph or explicit path.

Pathway Algorithm

Pathway explores candidate continuations and chooses a full routing decision for the requested destination or service. The main tuning question is how much search budget it needs before it reliably finds good candidates. It is also one of the in-tree routing engines that supports bounded deferred delivery of payloads: when a route enters partition mode, payloads can be retained through the shared retention boundary and later replayed on recovery or handoff.

Field Algorithm

Field maintains a continuously updated `field` model, searches over frozen snapshots, and publishes one corridor-style routing claim while allowing the concrete realization to move inside that corridor. It has an explicit bootstrap phase where weaker corridor claims can be published when evidence is coherent but not yet strong enough for steady admission. The tuning surface is now split deliberately: search breadth still lives in `FieldSearchConfig`, while continuation, promotion, continuity, and evidence behavior sit behind a separate internal operational policy surface that the experiments can sweep more cleanly. Field does carry forward bounded routing and service evidence, but that carry-forward supports corridor and service continuity rather than acting as a general payload store-and-forward cache.

Recommendation Logic

The recommendation score rewards settings that activate routes reliably, maintain route presence, tolerate harder stress levels, and (for the distance-vector engines) maintain stability. It penalizes route churn, maintenance failures, lost reachability, and prolonged degradation.

Profile-specific recommendations allow different operational priorities. Conservative profiles weight stability and failure avoidance more heavily, while aggressive or service-heavy profiles tolerate more risk.

`field` is calibrated on two surfaces in Part I: the generic route-visible recommendation surface and a separate regime-specific surface that explicitly scores corridor continuity, bootstrap upgrade quality, service continuity, and transition health.

When several nearby settings score about the same, the report prefers the middle of the acceptable range.

Reference Material

Appendix A contains the detailed transition, failure-boundary, profile, and `field`-specific calibration tables that support the main tuning recommendation.

Part II. Analysis

Reading Guide

Part II uses one repeated figure grammar across the engine sections:

- panels are scenario families
- the x-axis is the tuned control surface for that engine section
- solid lines with circle markers show the primary outcome view
- dashed lines with square markers show the cost, fragility, or startup view
- route-visible presence and activation are displayed as percentages

This makes the paired figures easier to compare: the first figure in a section shows how well the engine performs, and the second shows the cost, fragility, or control-motion price associated with that behavior.

3. BATMAN Classic Analysis

Findings

Recommended configuration: `batman-classic-2-1` (score=5071.3, activation=1000.0 permille, route presence=639.3 permille, max sustained stress=48).

The BATMAN Classic transition families separate most clearly at `batman-classic-10-5` (stability-total 34200, route presence 1000 permille).

BATMAN Classic's echo-only bidirectionality makes it consistently slower to materialize routes than BATMAN Bellman.

Transition Pressure Analysis

These two plots form an analytical pair: the first shows where stability accumulates across the transition families, and the second shows when those same settings first lose a route. The classic engine converges more slowly than other BATMAN variants due to its echo-only bidirectionality requirement.

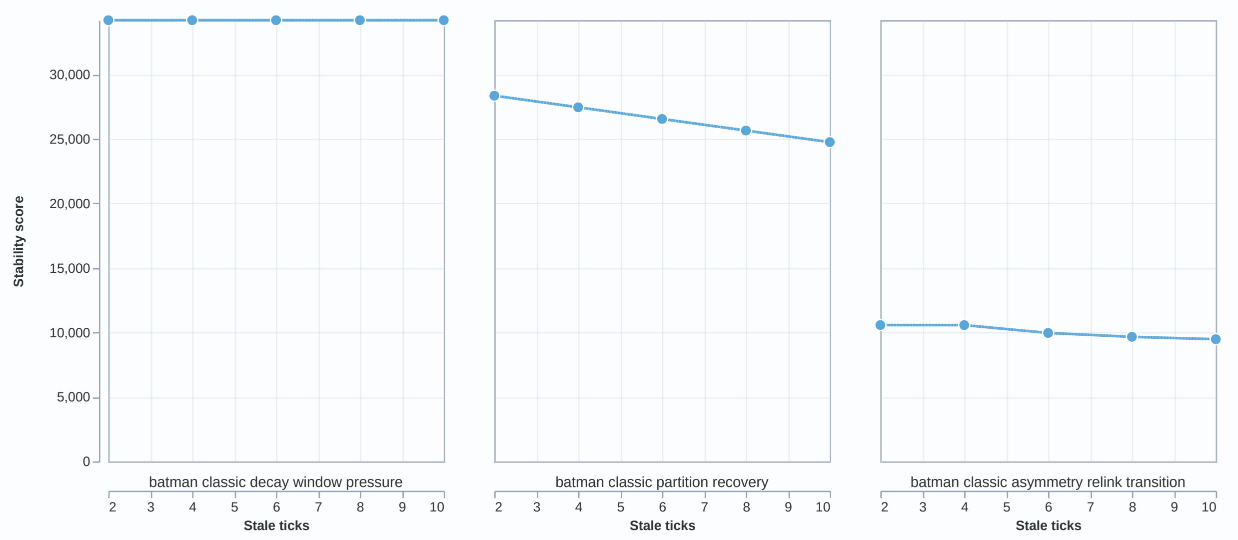

Figure 1. BATMAN Classic stability across transition families. Higher values are better: they indicate more sustained route quality across the scenario. Flatter high lines indicate a decay setting that stays robust as transition stress changes.

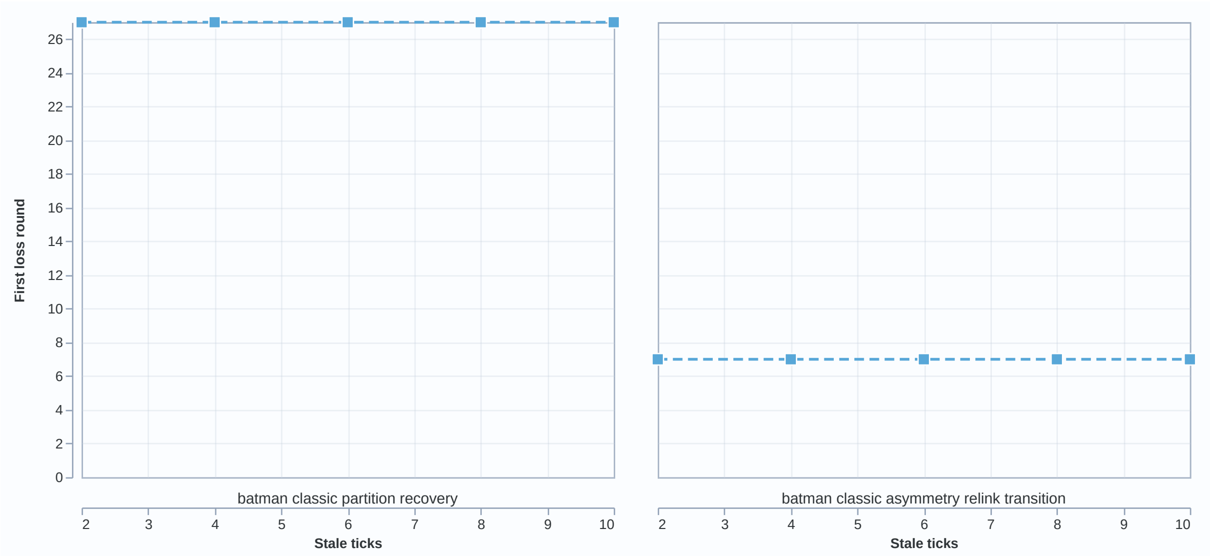

Figure 2. BATMAN Classic loss timing across transition families. Higher values mean the first route loss happens later, which is usually better. Sharp drops indicate settings that become brittle under the corresponding transition family.

4. BATMAN Bellman Analysis

Findings

Recommended configuration: `batman-bellman-1-1` (score=4289.2, activation=1000.0 permille, route presence=780.0 permille, max sustained stress=56).

The BATMAN Bellman transition families separate most clearly at `batman-bellman-1-1` (stability-total 6546, route presence 633 permille).

Severe asymmetric bridge loss remains a breakdown regime across the tested window range.

Transition Pressure Analysis

These two plots form an analytical pair: the first shows where stability accumulates across the transition families, and the second shows when those same settings first lose a route.

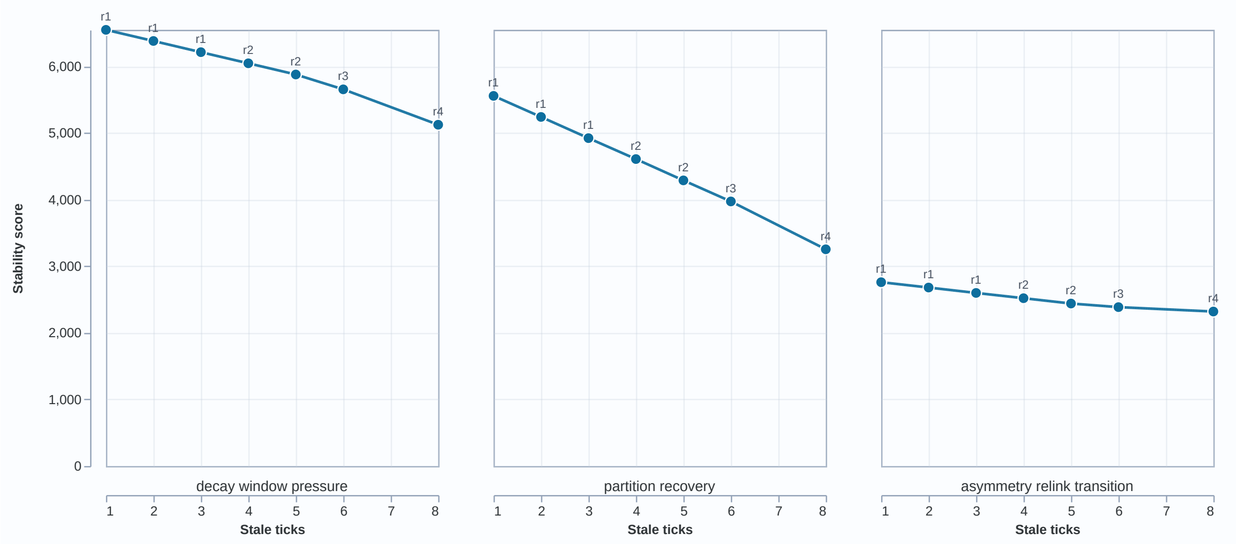

Figure 3. BATMAN Bellman stability across transition families. Higher values are better: they indicate more sustained route quality across the scenario. A broad plateau implies the stale-window setting is forgiving rather than narrowly tuned.

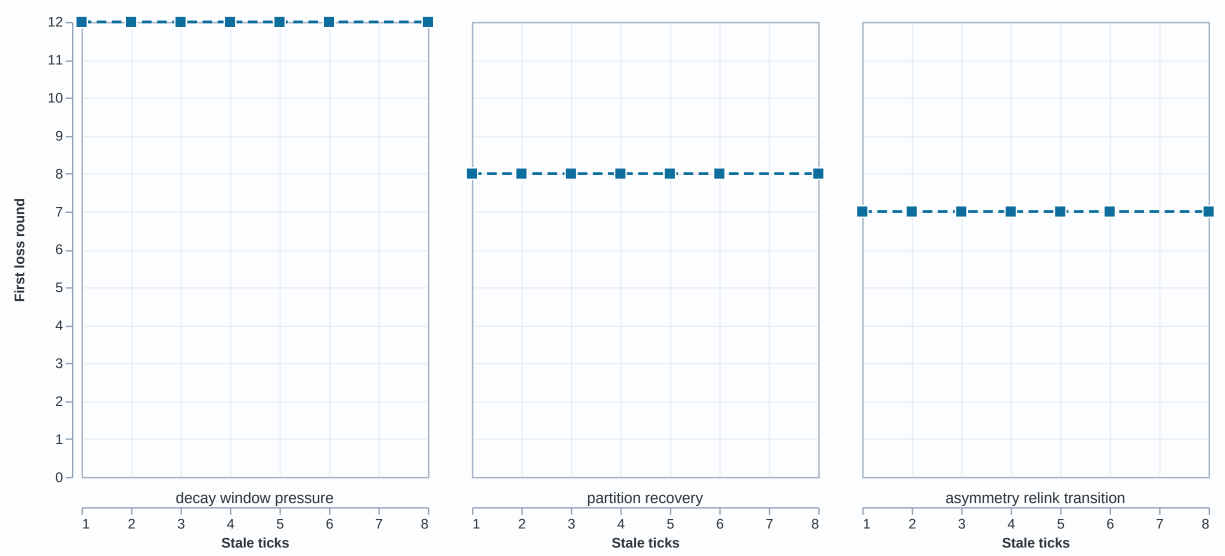

Figure 4. BATMAN Bellman loss timing across transition families. Higher values mean route loss is delayed further into the scenario, which is better. Early collapses indicate settings that cannot ride through the corresponding transition stress.

5. Babel Analysis

Findings

Recommended configuration: `babel-8-4` (score=4518.2, activation=1000.0 permille, route presence=922.3 permille, max sustained stress=52).

The Babel families separate most clearly at `babel-8-4` (stability-total 11798, route presence 1000 permille).

The feasibility distance table bounds convergence after partition recovery. Routes with the same seqno as pre-partition are infeasible until the next seqno increment.

Decay Window And Feasibility Analysis

The asymmetry-cost-penalty family is the primary differentiator: Babel's ETX formula penalizes poor reverse delivery more heavily than TQ, producing different route selection under identical topology. The partition-feasibility-recovery family shows the FD table's bounded infeasible-fallback window after partition clears.

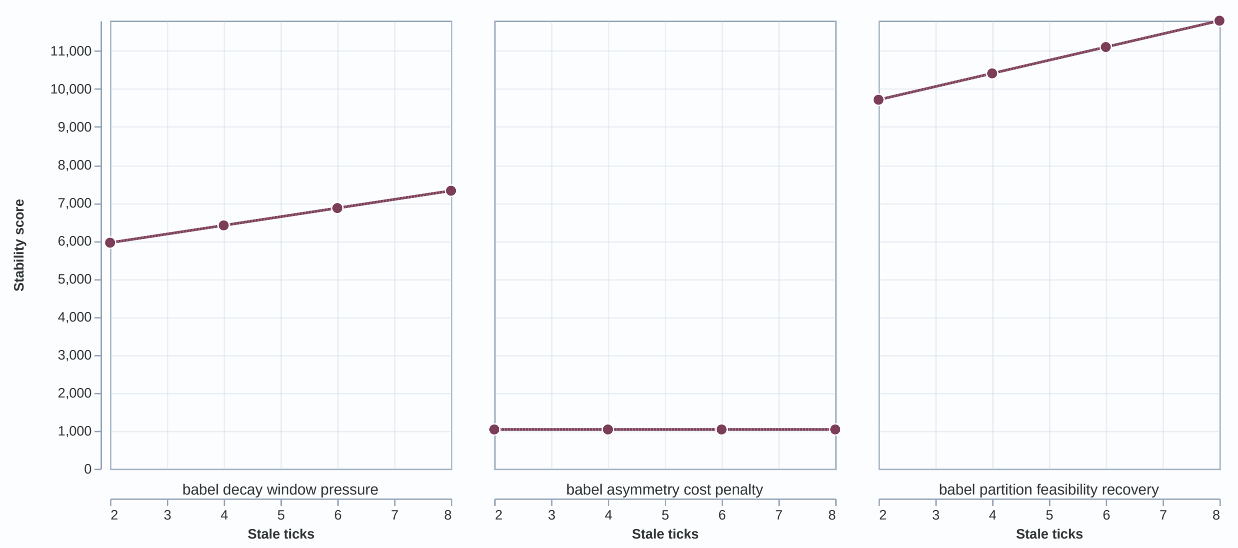

Figure 5. Babel stability across decay families. Higher values are better: they indicate more sustained feasible routing across the family. If the curve stays high as stale ticks increase, Babel is not very sensitive to the decay setting there.

6. OLSRv2 Analysis

Findings

Recommended configuration: `olsrv2-8-4` (score=5072.7, activation=1000.0 permille, route presence=792.2 permille, max sustained stress=52).

The OLSRv2 families separate most clearly at `olsrv2-2-1` (stability-total 23958, route presence 1000 permille).

The remaining stress point is asymmetric relink timing: full-topology knowledge helps after relays settle, but stale symmetric-link evidence can still leave one churn window where the best next hop is temporarily absent.

Topology Propagation And Churn Analysis

These plots answer the missing proactive link-state question directly: how quickly the full-topology engine stabilizes when links degrade, partitions clear, and relay roles shift. The topology-propagation and MPR-flooding families expose whether fresher HELLO and TC retention buys cleaner recovery or just unnecessary churn.

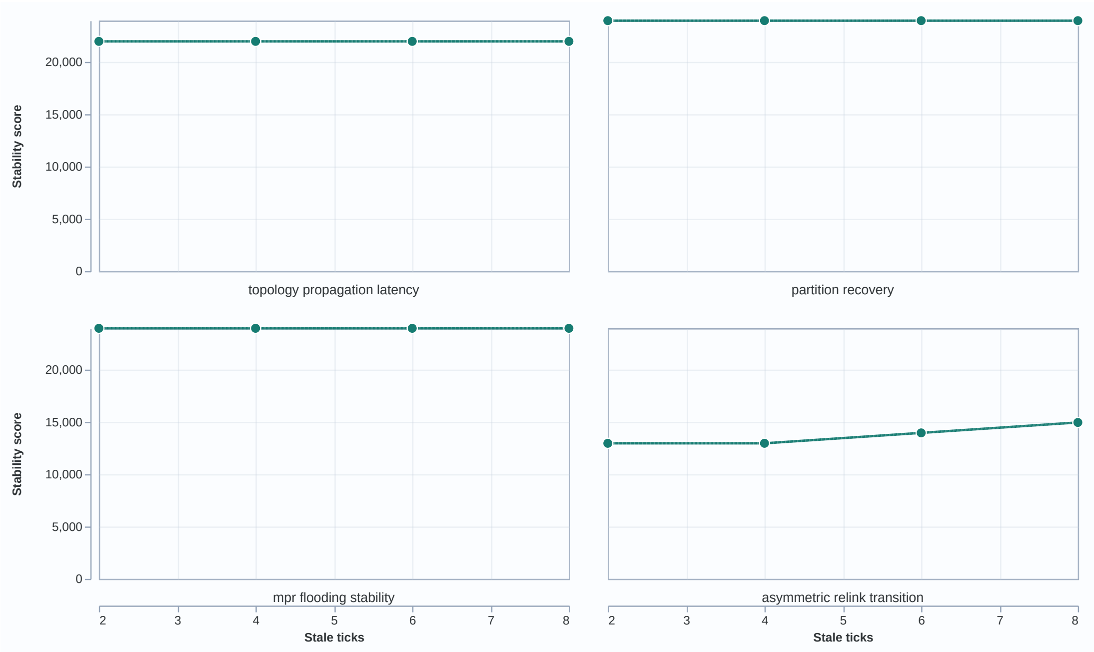

Figure 7. OLSRv2 stability across topology and churn families. Higher values are better: they indicate more sustained route quality through churn and relink events. A high flat region suggests the control-state lifetime is long enough to cover one topology-change cycle without overfitting.

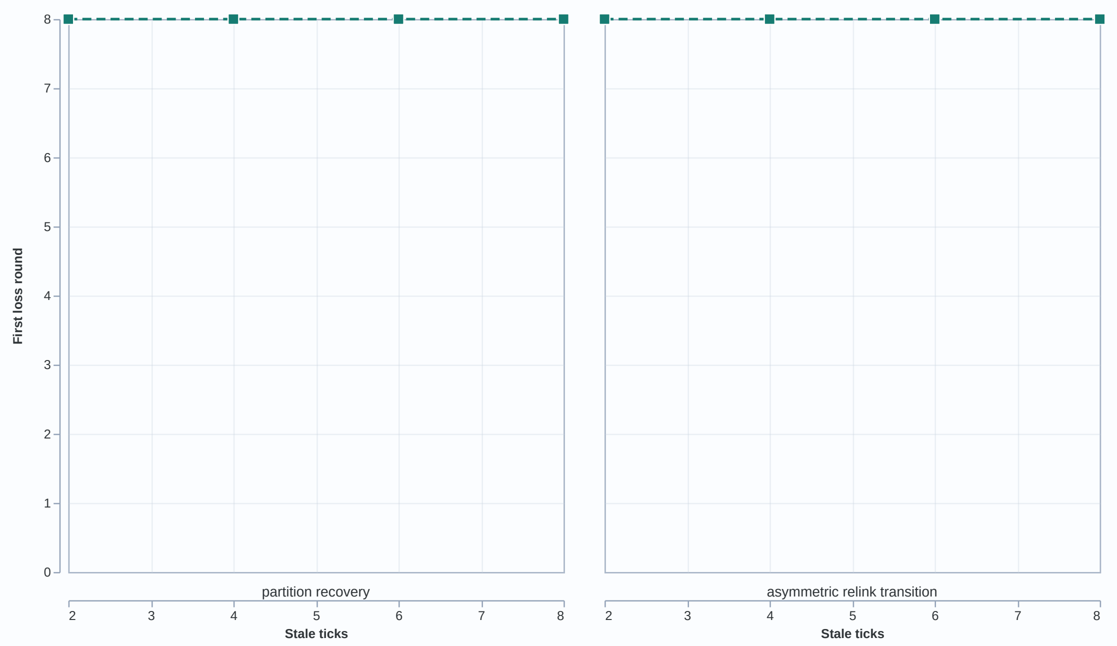

Figure 8. OLSRv2 loss timing across topology and churn families. Higher values mean the first loss occurs later, which is better. Lower or falling values indicate churn windows where stale symmetric-link or TC state leaves the engine exposed.

7. Scatter Analysis

Findings

Recommended configuration: `scatter-balanced` (score=5292.4, activation=1000.0 permille, route presence=900.5 permille, max sustained stress=90).

Scatter separates most clearly in `scatter-low-rate-transfer-threshold`, where the owner-side runtime surface records handoff 14, constrained occupancy 16, and bridging 0 while normalized route presence stays at 888 permille.

The key Scatter contrast is architectural rather than path-optimal: it is the opaque, partition-tolerant baseline that keeps payload custody local and bounded instead of searching a full path or publishing a corridor envelope.

These two figures put `scatter` on the same tuning-sweep footing as the other engines without pretending that route presence is the whole story. The first figure shows the route-visible tie, and the second shows the threshold-runtime behavior that now separates the maintained `balanced`, `conservative`, and `degraded-network` profiles.

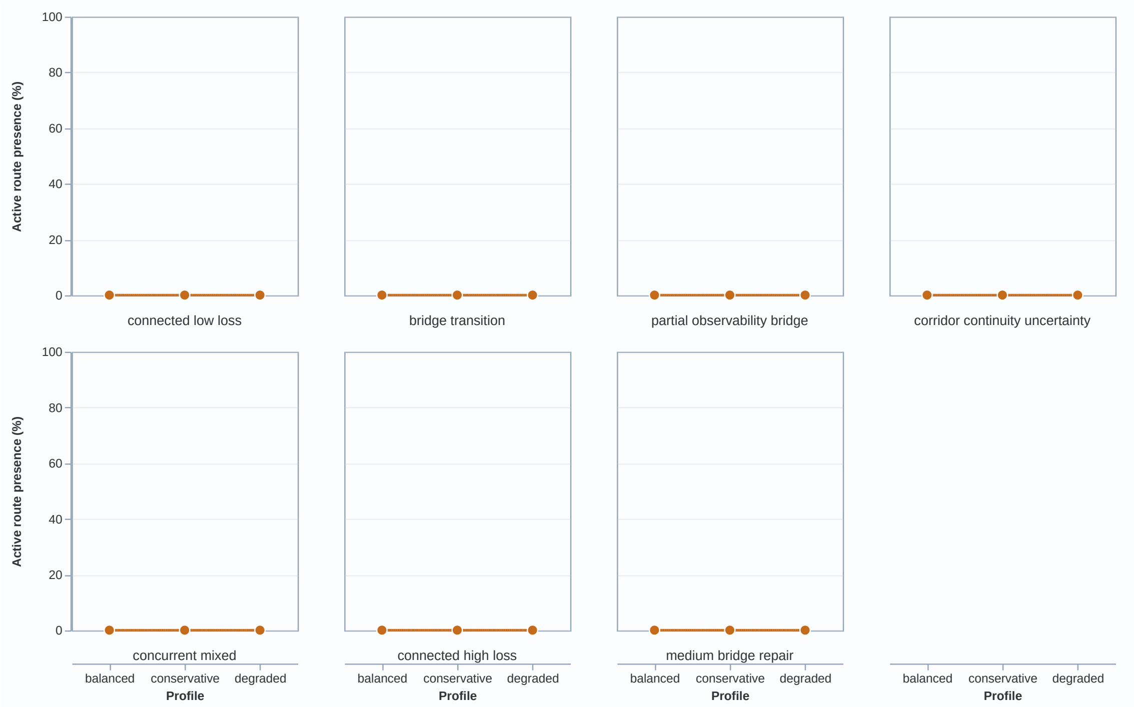

Figure 9. Scatter active route presence by maintained profile. Higher values are better because the route stays present for more of the objective-active window. Each panel uses the same profile order so the `scatter` sweep can be compared directly across families.

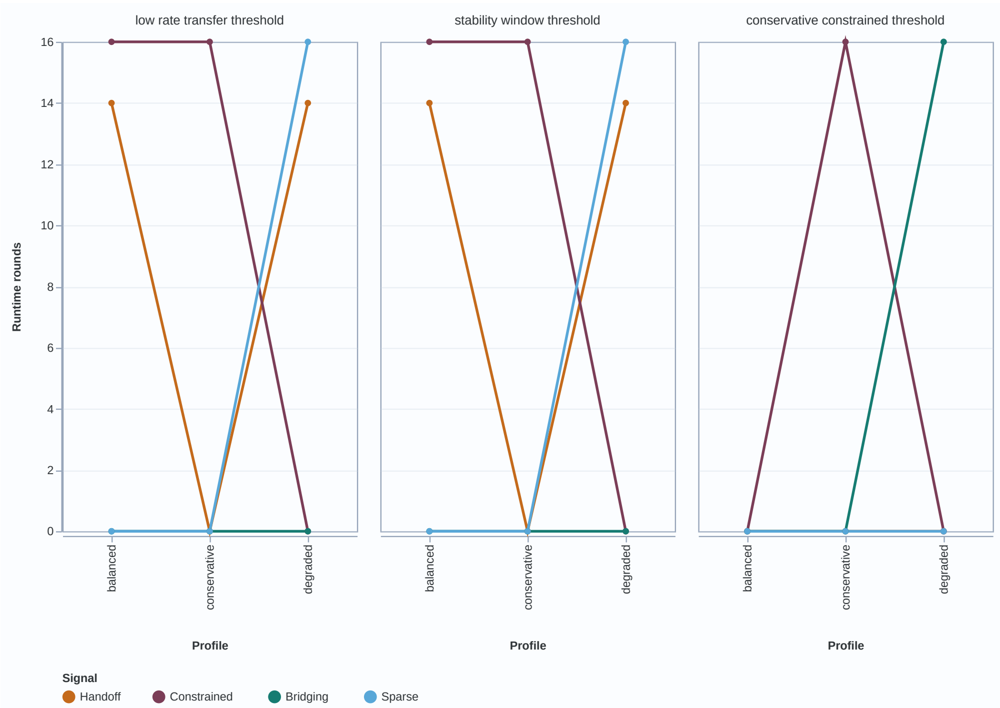

Figure 10. Scatter threshold-runtime behavior by maintained profile. Higher values are not universally better: the point is to show which profiles spend rounds in handoff, constrained, bridging, or sparse regimes under the new threshold families. This is the informative Scatter tuning surface when route-visible outcomes tie.

8. Pathway Analysis

Findings

Recommended configuration: `pathway-2-hop-lower-bound` (score=4166.3, activation=1000.0 permille, route presence=876.3 permille, max sustained stress=58).

Pathway budget 1 is the cliff edge: activation=1000 permille.

Budgets at and above `pathway-2-hop-lower-bound` form the stable floor.

Budget Figures

These two figures show the budget question from two angles: how much route-visible outcome extra budget buys, and what startup or fragility cost remains.

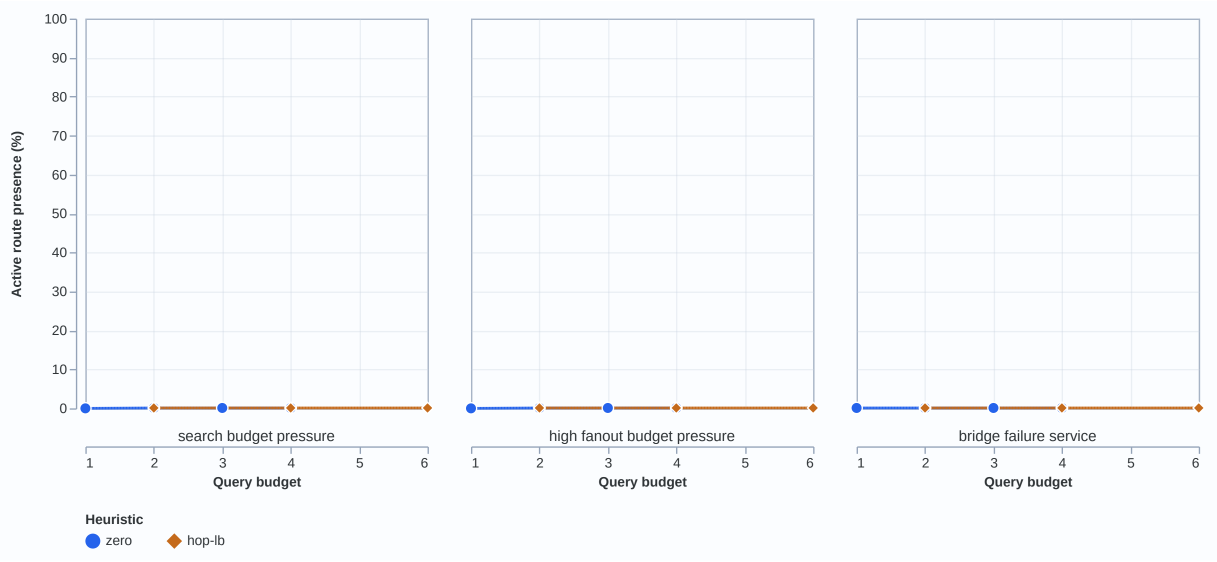

Figure 11. Pathway active route presence by search budget. Higher values are better: they indicate the route is present for more of the objective-active window. The y-axis is shown as a percentage so the budget sweep can be read directly against the other route-visible outcome figures in Part II.

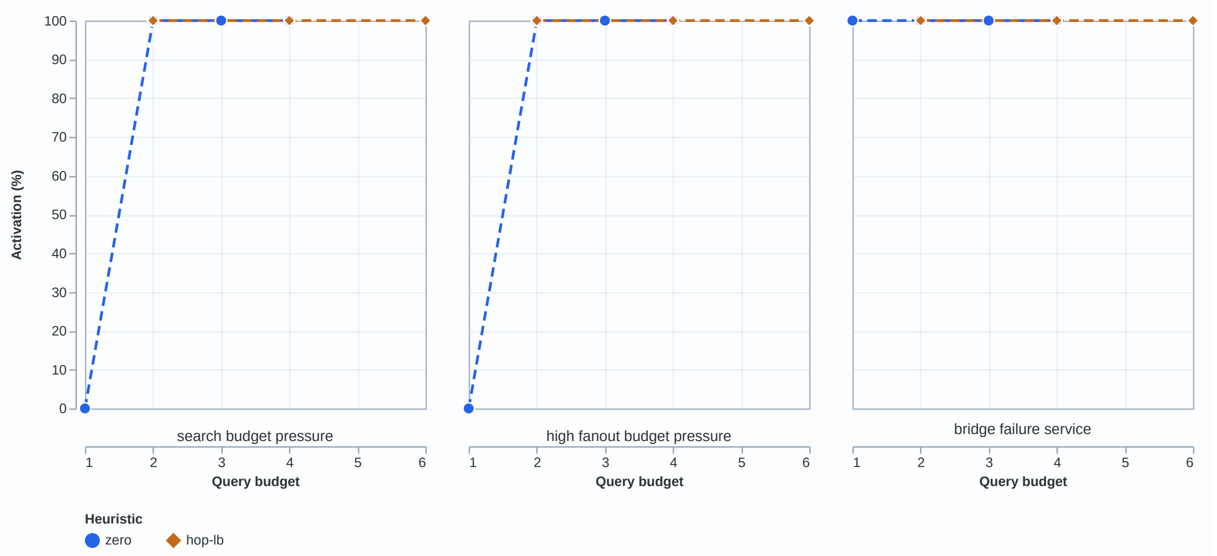

Figure 12. Pathway activation by search budget. Higher values are better: they indicate objectives activate successfully more often. The y-axis is shown as a percentage, and step changes reveal the budget threshold where Pathway moves from under-search to reliable activation.

9. Field Analysis

Findings

Recommended configuration: `field-4-zero-p1-f140-n180` (score=4606.5, activation=1000.0 permille, route presence=780.9 permille, max sustained stress=60).

Field separates where corridor continuity and reconfiguration cost both matter. `field-6-hop-lower-bound-p2-f130-n130` keeps route presence at 1000 permille while holding continuation shifts to 0.

Corridor-continuity profile: bootstrap activation 444.4 permille, hold 666.7 permille, narrow 0.0 permille, upgrade 222.2 permille, withdrawal 555.6 permille, degraded-steady occupancy 69.4 permille, service carry-forward 6555.6 permille, asymmetric shift success 111.1 permille. Dominant commitment resolution `Pending`, last recovery outcome `ContinuationRetained`, continuity band `Steady`, continuity transition `EnteredDegradedSteady`, last decision `Withdraw`, blocker `SupportTrend`.

Route presence is close across the tested range. The service-oriented knobs separate configs in continuation-shift count, service carry-forward, and narrowing behavior.

The maintained families produce router-visible activation and route presence, with the bootstrap phase directly visible in replay and recovery surfaces.

The asymmetric-envelope and bridge anti-entropy families test corridor continuity under realization movement. The partial-observability and bootstrap-upgrade families test bootstrap promotion and withdrawal. The service-overlap, freshness-inversion, and publication-pressure families test service-corridor continuity under broader publication, stronger freshness weighting, or earlier narrowing.

The service-corridor publication and materialization path is the key enabler for Field's route-visible behavior. The tuning question is which continuity style is preferable: narrower lower-churn publication or broader reselection with more carried-forward optionality.

Corridor Figures

The first figure shows route-visible continuity across corridor-oriented families. The second shows the control-motion cost paid to preserve that continuity. Together they distinguish a healthy corridor default from an unstable bootstrap regime.

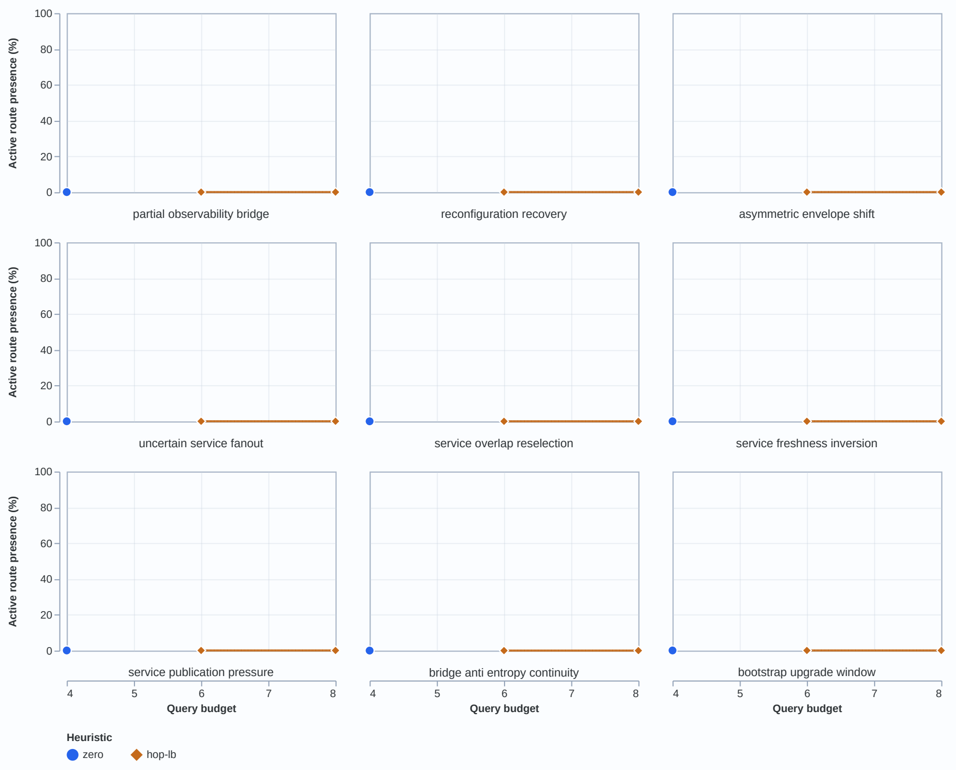

Figure 13. Field active route presence by search budget. Higher values are better: they indicate the admitted corridor stays available for more of the active window. The y-axis is shown as a percentage so the continuity outcome can be compared directly against the other Part II route-visible figures.

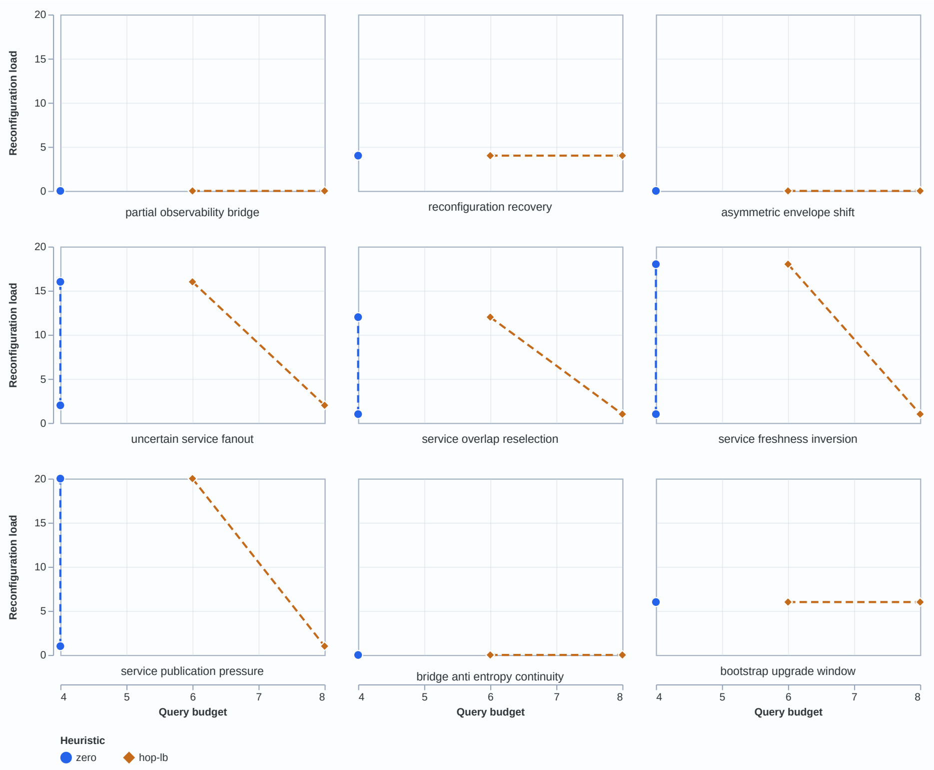

Figure 14. Field corridor reconfiguration by search budget. Lower values are generally better because they indicate less continuation churn and fewer search-driven reconfigurations. Rising lines mean the engine is paying more control-motion cost to maintain continuity.

10. Comparative Analysis

Mixed-Engine Comparison

`comparison-bridge-transition`: selected_rounds_leader=`batman-classic`, activation=1000, active_route_presence=1000

`comparison-concurrent-mixed`: selected_rounds_leader=`pathway`, activation=1000, active_route_presence=1000

`comparison-connected-high-loss`: selected_rounds_leader=`pathway`, activation=1000, active_route_presence=1000

`comparison-connected-low-loss`: selected_rounds_leader=`scatter`, activation=1000, active_route_presence=1000

`comparison-corridor-continuity-uncertainty`: selected_rounds_leader=`pathway`, activation=1000, active_route_presence=1000

`comparison-large-core-periphery-high`: selected_rounds_leader=`pathway`, activation=1000, active_route_presence=1000

`comparison-large-core-periphery-moderate`: selected_rounds_leader=`pathway`, activation=1000, active_route_presence=1000

`comparison-large-multi-bottleneck-high`: selected_rounds_leader=`pathway`, activation=1000, active_route_presence=1000

`comparison-large-multi-bottleneck-moderate`: selected_rounds_leader=`pathway`, activation=1000, active_route_presence=1000

`comparison-medium-bridge-repair`: selected_rounds_leader=`pathway`, activation=1000, active_route_presence=1000

`comparison-multi-flow-asymmetric-demand`: selected_rounds_leader=tie, activation=1000, active_route_presence=1000

`comparison-multi-flow-detour-choice`: selected_rounds_leader=tie, activation=1000, active_route_presence=1000

`comparison-multi-flow-shared-corridor`: selected_rounds_leader=tie, activation=1000, active_route_presence=1000

`comparison-partial-observability-bridge`: selected_rounds_leader=`olsrv2`, activation=1000, active_route_presence=1000

`comparison-pathway-budget-boundary`: selected_rounds_leader=`pathway`, activation=1000, active_route_presence=1000

`comparison-stale-asymmetric-region`: selected_rounds_leader=`pathway`, activation=1000, active_route_presence=1000

`comparison-stale-observation-delay`: selected_rounds_leader=`pathway`, activation=1000, active_route_presence=1000

`comparison-stale-recovery-window`: selected_rounds_leader=`pathway`, activation=1000, active_route_presence=1000

Head-To-Head Engine Sets

These rows show what each stack does when it is the only available routing surface for that host set. This is a fixed representative-profile benchmark surface, not the calibrated-best profile surface from Part I.

`head-to-head-bridge-transition`: best engine set=`pathway`, activation=1000 permille, active route presence=1000 permille.

`head-to-head-concurrent-mixed`: best engine set=`pathway`, activation=1000 permille, active route presence=1000 permille.

`head-to-head-connected-high-loss`: best engine set=`pathway`, activation=1000 permille, active route presence=1000 permille.

`head-to-head-connected-low-loss`: best engine set=`batman-classic`, activation=1000 permille, active route presence=1000 permille.

`head-to-head-corridor-continuity-uncertainty`: best engine set=`pathway`, activation=1000 permille, active route presence=1000 permille.

`head-to-head-large-core-periphery-high`: best engine set=`batman-bellman`, activation=1000 permille, active route presence=1000 permille.

`head-to-head-large-core-periphery-moderate`: best engine set=`batman-bellman`, activation=1000 permille, active route presence=1000 permille.

`head-to-head-large-multi-bottleneck-high`: best engine set=`batman-bellman`, activation=1000 permille, active route presence=1000 permille.

`head-to-head-large-multi-bottleneck-moderate`: best engine set=`batman-bellman`, activation=1000 permille, active route presence=1000 permille.

`head-to-head-medium-bridge-repair`: best engine set=`batman-bellman`, activation=1000 permille, active route presence=1000 permille.

`head-to-head-multi-flow-asymmetric-demand`: best engine set=`batman-bellman`, activation=1000 permille, active route presence=1000 permille.

`head-to-head-multi-flow-detour-choice`: best engine set=`batman-bellman`, activation=1000 permille, active route presence=1000 permille.

`head-to-head-multi-flow-shared-corridor`: best engine set=`batman-bellman`, activation=1000 permille, active route presence=1000 permille.

`head-to-head-partial-observability-bridge`: best engine set=`batman-classic`, activation=1000 permille, active route presence=1000 permille.

`head-to-head-stale-asymmetric-region`: best engine set=`batman-bellman`, activation=1000 permille, active route presence=1000 permille.

`head-to-head-stale-observation-delay`: best engine set=`batman-bellman`, activation=1000 permille, active route presence=1000 permille.

`head-to-head-stale-recovery-window`: best engine set=`batman-bellman`, activation=1000 permille, active route presence=1000 permille.

Head-To-Head Regimes

The head-to-head regimes are:

- `connected-low-loss`: easy connected route where all engines should establish a route.
- `connected-high-loss`: repairable connected route over a lossy bridge.
- `bridge-transition`: bridge that degrades, partitions, and restores.
- `medium-bridge-repair`: moderate bridge degradation with a repair window rewarding durable recovery without needing a fully mixed workload.
- `partial-observability-bridge`: bridge case with Field bootstrap summaries for corridor-style routing under incomplete evidence.

- `corridor-continuity-uncertainty`: intermittent degradation and restoration rewarding corridor continuity.
- `concurrent-mixed`: multiple active objectives testing mixed-workload behavior.

Limitations And Next Steps

These recommendations are only as good as the simulated regime corpus. A flat curve can mean genuine robustness or that the sweep has not found the most informative failure boundary.

The BATMAN Bellman corpus exposes recoverable transition differences, but asymmetry-plus-bridge families remain hard failures. The BATMAN Classic corpus confirms slower convergence and tight clustering of decay window settings. The Babel corpus shows measurably different behavior under asymmetric conditions, with the FD table visible in partition recovery, but decay window settings do not yet separate sharply. The OLSRv2 corpus separates most clearly on topology propagation, MPR flooding stability, and asymmetric relink timing, but the maintained window sweep is still narrow enough that several settings remain tied on route visibility. The Pathway corpus identifies the minimum viable budget floor with a wide plateau above it.

The Field corpus reaches the route-visible boundary with an explicit bootstrap phase and working service-corridor path, but tested settings cluster tightly. The bridge anti-entropy and bootstrap-upgrade families allow the report to distinguish between underexercise and real weakness. The remaining limitation is that tested settings cluster closely, leaving room for more discriminating future regime design.

The remaining routing-fitness experiments narrow the route-visible decision to one explicit envelope rather than several open questions. Within the tested search-burden, shared-broker, and stale-observation bands, the hybrid search-plus-maintenance direction is supported as fit for purpose; pure search degrades earlier once diameter and shared bottlenecks both rise. The remaining limitation is now extrapolation beyond this maintained envelope, not uncertainty about the basic direction inside it.

The full mixed-engine and head-to-head tables are collected in Appendix B. The main body keeps the figures and takeaways.

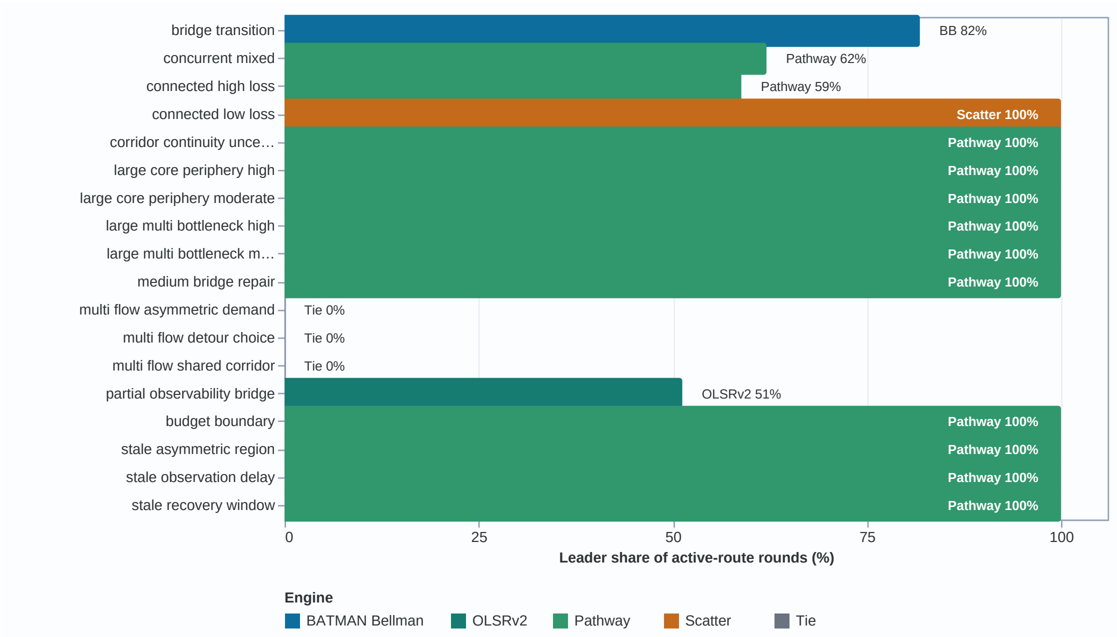

Figure 15. Mixed-engine router arbitration by comparison regime. Bar color marks the engine the deterministic router selected most often in the mixed stack, while the in-bar label shows that leader's share of active-route rounds. This is an arbitration view, not a standalone performance comparison: values near 100% mean the router effectively stuck with one engine for that regime, while lower percentages mean arbitration was more split.

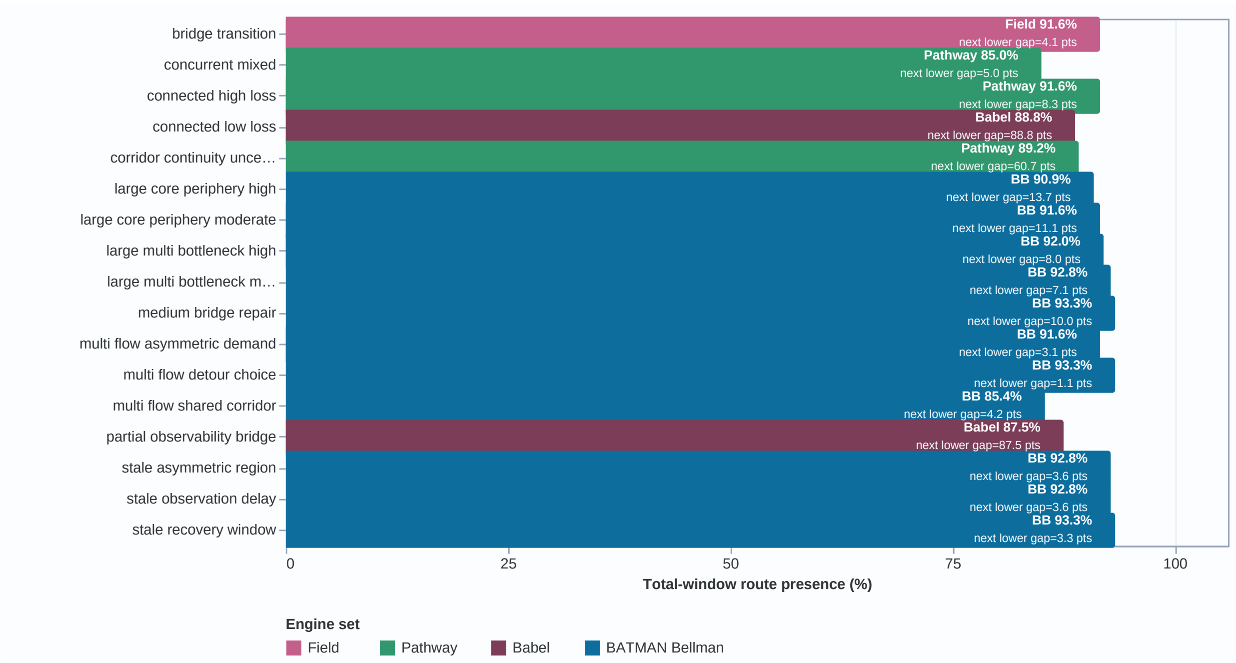

Figure 16. Head-to-head standalone capability by comparison regime. Longer bars are better: they mark the engine with the highest total-window route presence when run alone, and bar width encodes that route-presence level directly. This is the standalone capability view for the same regime families, so a small `next lower gap` means the engines cluster tightly at the top while a large gap means the scenario cleanly separates the leading tier from the rest.

Figure 17. Timing view for the same head-to-head families. The left panel shows who gets a route up first; the right panel shows who keeps that route longest before the first observed loss.

Figure 18. Robustness view for the current recommended defaults. This figure shows which engines combine high route presence with lower regime-to-regime spread, rather than only maximizing the mean.

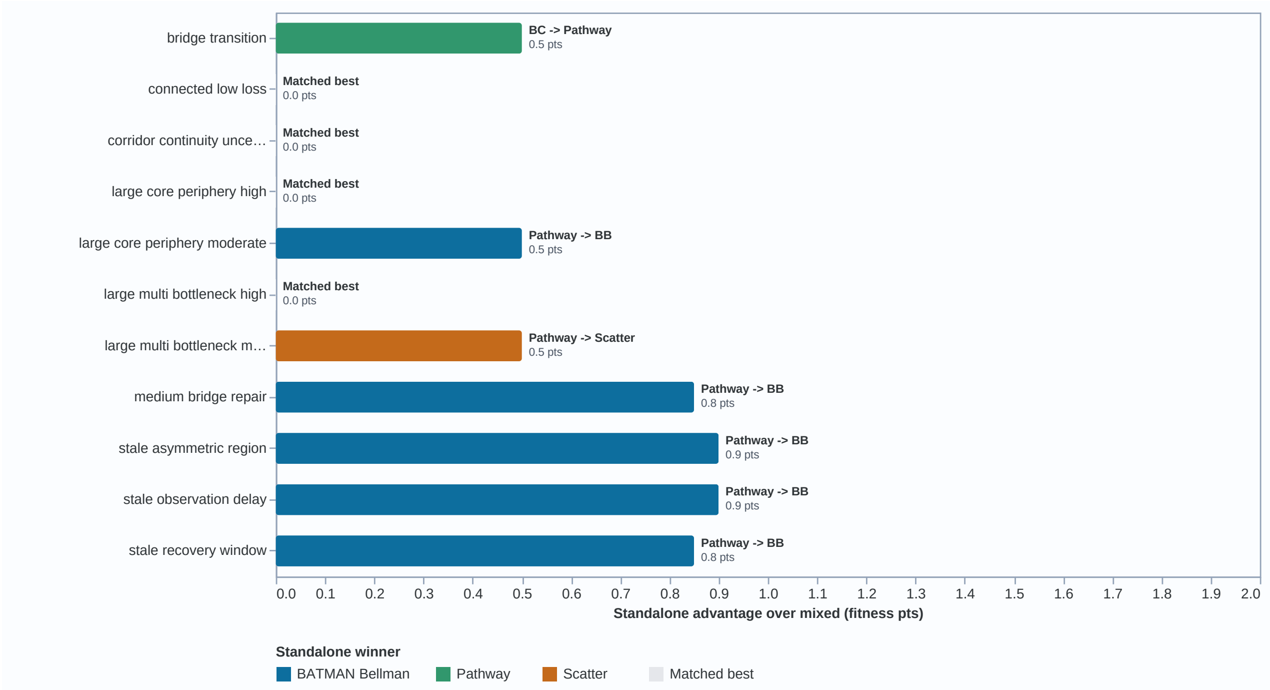

Figure 19. Direct fitness gap between what the mixed router achieved and what the best standalone engine would have achieved for the same named family. Longer bars are larger gaps after accounting for activation success, total-window route presence, materialization delay, route churn, and activation failures; bar color identifies the standalone winner defining that gap. This is the explicit bridge between the arbitration story in Figure 15 and the capability story in Figure 16.

- `connected-low-loss` and `partial-observability-bridge` are broad tie regimes; in the latter, `babel`, `batman-bellman`, `batman-classic`, `field`, `olsrv2`, `pathway`, `pathway-batman-bellman`, and `scatter` all reach 1000 permille active-window route presence.
- The hard route-visible bridge families are the clearest separators: `connected-high-loss` is led by `pathway`, `pathway-batman-bellman`, and `scatter` at 1000 permille, while `bridge-transition` is shared by `field`, `pathway`, and `scatter` at 1000 permille.
- `medium-bridge-repair` also stays broad at the top: `batman-bellman`, `field`, `pathway`, `pathway-batman-bellman`, and `scatter` all reach 1000 permille.
- Mixed workloads still favor explicit search: `concurrent-mixed` is led by `pathway`, `pathway-batman-bellman`, and `scatter` at 1000 permille.
- `field` stays competitive in the hard bridge families at 909 permille in `connected-high-loss` and 1000 permille in `bridge-transition`, but it drops to 320 permille in `corridor-continuity-uncertainty`.

Part II Takeaways

- The routing comparison does not collapse to one universal winner. In the mixed-engine matrix, `connected-low-loss` is led by `scatter`, `connected-high-loss` by `pathway`, and `concurrent-mixed` by `pathway`.
- Among the maintained proactive next-hop defaults, `babel` and `olsrv2` are the strongest contrasting baselines: `babel-8-4` captures the asymmetry-sensitive distance-vector case, while `olsrv2-8-4` is the full-topology baseline when HELLO and TC propagation have time to pay off.
- Explicit search still matters when the workload is mixed rather than purely hop-by-hop. In the head-to-head matrix, `concurrent-mixed` is led by `pathway`, `pathway-batman-bellman`, and `scatter` at 1000 permille route presence.

- The mixed router leaves large performance on the table in the hardest bridge families: `connected-high-loss` settles on `pathway` at 1000 permille while standalone `pathway`, `pathway-batman-bellman`, and `scatter` reach 1000, and `bridge-transition` settles on `batman-classic` at 1000 while standalone `field`, `pathway`, and `scatter` reach 1000.
- `field` is corridor-oriented rather than universal: it stays competitive in `connected-high-loss` at 909 permille and in `bridge-transition` at 1000 permille, but in `corridor-continuity-uncertainty` it drops to 320 permille while `pathway`, `pathway-batman-bellman`, and `scatter` lead at 1000.

11. Large-Population Findings

These additions extend the maintained corpus beyond the small connected and single-bridge families so the report can ask three larger-network questions directly: control-plane scaling under fanout and diameter, multi-bottleneck fragility under overlapping repair pressure, and diffusion phase transitions in larger clustered populations.

The route-visible track adds moderate and high large-pop bands for a mixed-density core-periphery family and a multi-bottleneck repair family. The diffusion track adds sparse-threshold, congestion-threshold, and regional-shift continuity families at moderate and high bands.

Large-Population Route Summary

Compact route-visible large-population surface by topology class and engine set.

Topology	Engine Set	Small	Moderate	High	dHigh	High Loss
Diameter / fanout scaling	BATMAN Classic	888	805	772	-116	-
Diameter / fanout scaling	BATMAN Bellman	888	916	909	21	-
Diameter / fanout scaling	Babel	888	805	772	-116	-
Diameter / fanout scaling	OLSRv2	888	805	772	-116	-
Diameter / fanout scaling	Pathway	888	916	909	21	-
Diameter / fanout scaling	Scatter	888	916	909	21	-
Diameter / fanout scaling	Pathway + BATMAN Bellman	888	916	909	21	-
Diameter / fanout scaling	Field	888	55	45	-843	6
Multi-bottleneck repair	BATMAN Classic	833	857	840	7	-
Multi-bottleneck repair	BATMAN Bellman	933	928	920	-13	-
Multi-bottleneck repair	Babel	833	857	840	7	-
Multi-bottleneck repair	OLSRv2	833	857	840	7	-
Multi-bottleneck repair	Pathway	933	928	920	-13	-
Multi-bottleneck repair	Scatter	933	928	920	-13	-
Multi-bottleneck repair	Pathway + BATMAN Bellman	933	928	920	-13	-
Multi-bottleneck repair	Field	933	285	200	-733	14

Table 2. Topology is the analytical question family. Engine Set is the standalone routing stack. Small, Moderate, and High are total-window route-presence means for the maintained size bands. `dHigh` is the high-band minus small-band route-presence delta. High Loss is the mean first-loss round in the high-band family.

Large-Population Diffusion Transitions

Representative collapse, viable, and explosive profiles for each maintained large-population diffusion family.

Question	Size	Collapse	Viable	Explosive
Congestion threshold	Moderate	Field congestion s2	Transition broad	-
Congestion threshold	High	Transition broad	-	-
Regional shift continuity	Moderate	Transition tight	Field congestion	Field
Regional shift continuity	High	Transition tight	Field scarcity s3	Field congestion
Sparse threshold	Moderate	Transition tight	Field privacy s4	Field continuity
Sparse threshold	High	Transition tight	Field privacy	Field congestion s3

Table 3. Question is the diffusion question family. Size is the maintained population band. Collapse / Viable / Explosive show the representative configuration for that boundedness state when one was observed.

Large-Population Figure Context

These figures separate route-visible scaling, route-visible fragility, and diffusion transition behavior across the maintained large-population corpus.

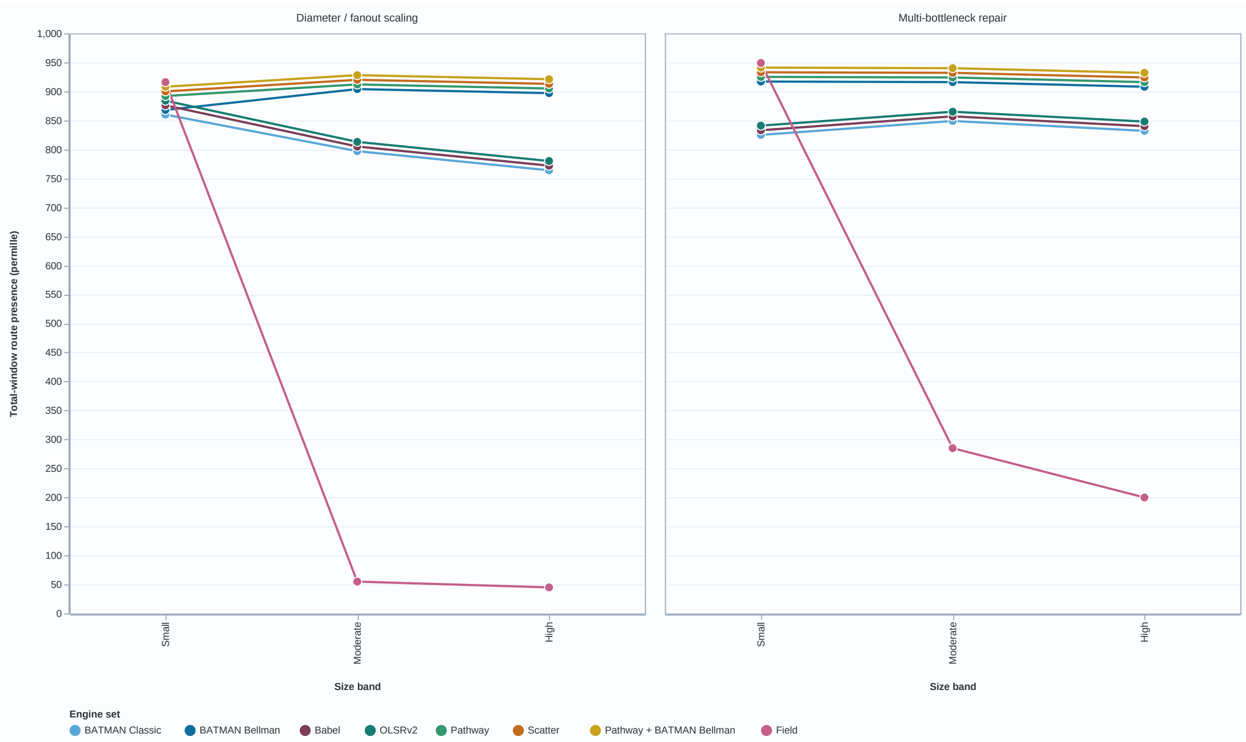

Figure 22. Route-visible performance by size band for the maintained large-population families. Each panel fixes one topology class and traces how each standalone engine set moves from the small baseline into the moderate and high bands. Read the slope more than the starting point: flatter lines mean the engine keeps its route-visible behavior as the population grows, while steep drops mean scaling pressure is exposing a control or search limit. An engine that stays high across all three bands is scaling cleanly in that topology class rather than winning only in the smallest case.

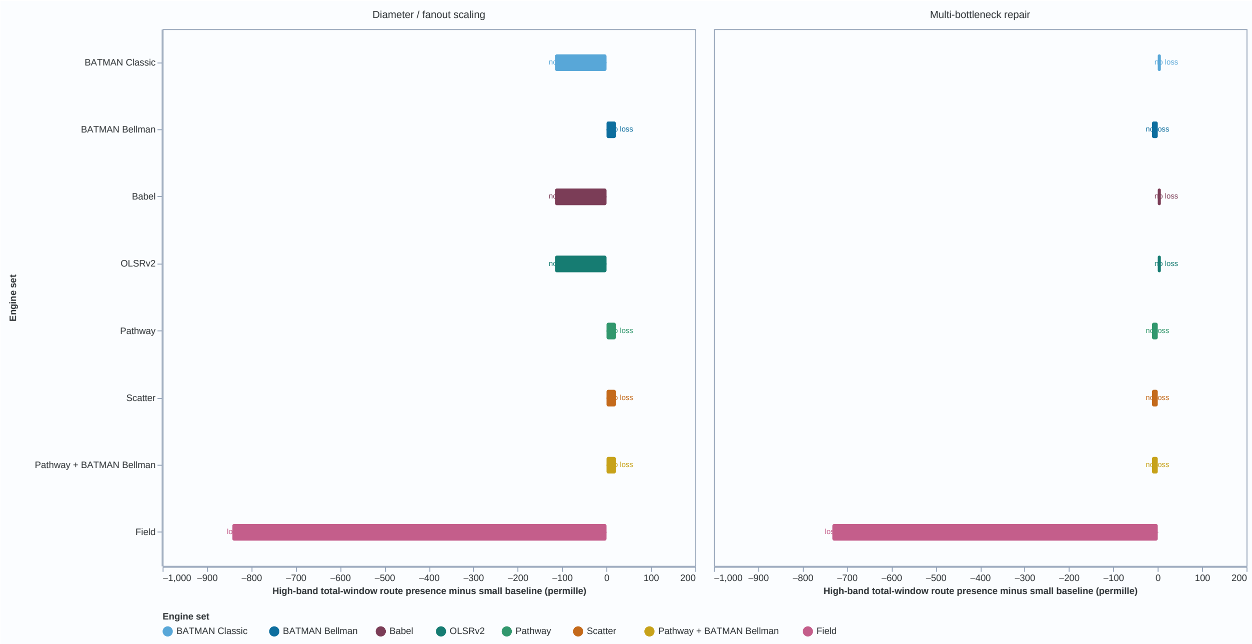

Figure 23. Small-to-high route-presence drop for the maintained large-population route-visible classes. More negative bars mean stronger degradation as the graph grows or bottlenecks multiply. The inline loss label gives the mean first-loss round in the high band. Use this as a degradation summary rather than a raw performance chart: bars near zero mean the engine preserved most of its small-band behavior, while large negative bars mean the larger graph is causing material fragility. Earlier high-band loss labels indicate that the engine is not only degrading more, but also failing sooner under the larger-population regime.

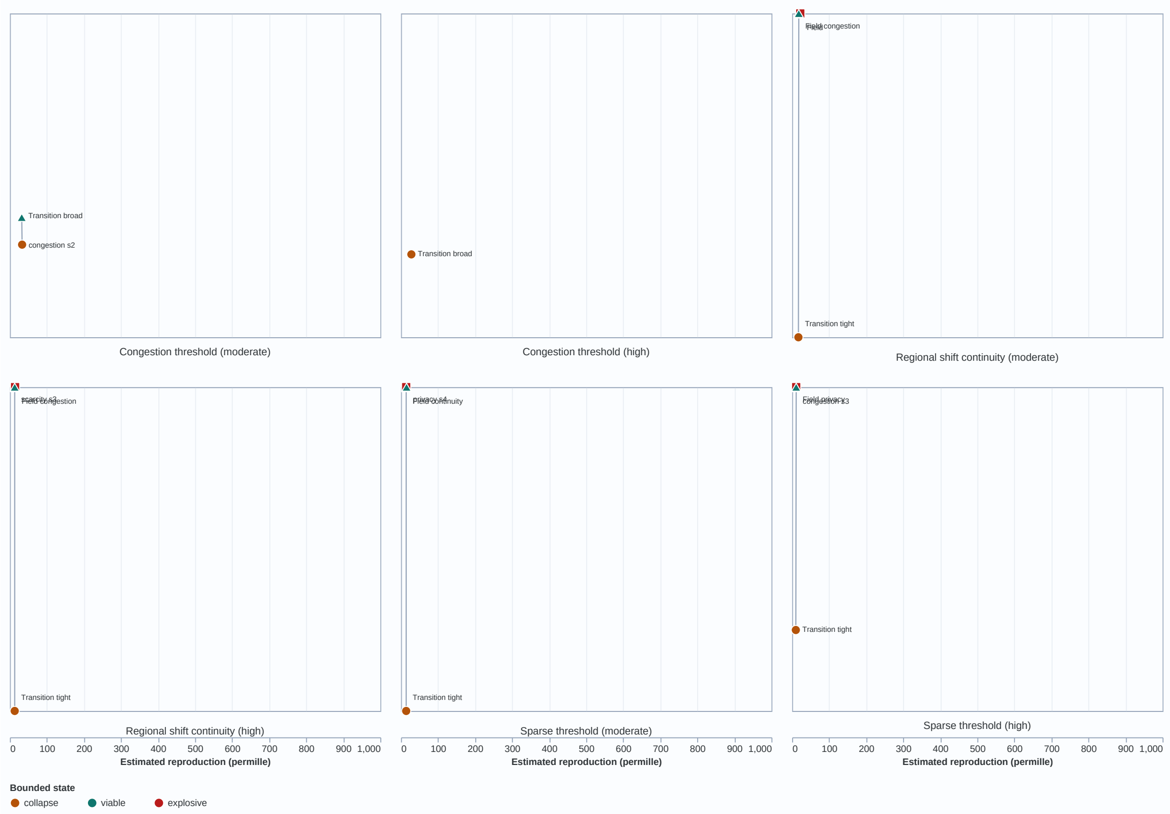

Figure 24. Representative bounded-state points for the maintained large-population diffusion families. Each panel shows delivery versus estimated reproduction for the best observed collapse, viable, and explosive representatives in that family, making the transition surface visible without scanning the full raw matrix. Interpret the point positions by quadrant: the useful region is high delivery with bounded reproduction, while low-delivery collapse points and high-R explosive points show the two ways an engine can fail at scale. Engines that keep their viable representative well separated from both failure modes are handling larger-population diffusion pressure more cleanly.

Large-Population Takeaways

- In the high large-pop route-visible bands, the core-periphery family is shared by `batman-bellman`, `pathway`, `pathway-batman-bellman`, and `scatter` at 909 permille, while the high multi-bottleneck family is shared by `batman-bellman`, `pathway`, `pathway-batman-bellman`, and `scatter` at 920 permille.
- The steepest diameter / fanout drop is `field` at -843 permille from the small baseline to the high band, and the steepest multi-bottleneck drop is `field` at -733 permille.
- `scatter` and `field` are the clearest route-visible large-population losers: in the high core-periphery band they fall to 909 and 45 permille, and in the high multi-bottleneck band they reach only 920 and 200 permille.
- The combined `pathway-batman-bellman` stack materially closes the multi-bottleneck gap versus plain `pathway`, reaching 920 permille in the high band versus 920 for `pathway` alone.
- The sparse-threshold high band still shows viable `field-privacy` against explosive `field-congestion-search-3`, the congestion-threshold moderate band separates viable `transition-broad` from collapse `field-congestion-search-2`, the congestion-threshold high band is currently only `collapse`, and the regional-shift high band still spans `collapse, explosive, viable`.

12. Routing-Fitness Remaining Questions

The earlier route-visible and large-population sections were enough to choose a design direction, but not enough to close the remaining fitness-for-purpose questions. Three gaps remained: where explicit search stops being sufficient by itself, what shared-broker contention does to the weakest flow rather than the mean, and whether delayed or asymmetric observations cause the router to cling to dead routes after the ground truth changes.

This section answers those questions directly. The crossover sweep reuses the maintained large-pop route-visible bands as a controlled search-burden versus maintenance-benefit surface. The multi-flow families force several simultaneous objectives through shared brokers. The stale-repair families inject host-specific lag windows so different regions act on delayed topology knowledge even though the underlying ground truth changes on one deterministic schedule.

Across these summaries, route churn is the report-primary cost metric and the route-observation count is retained as a control-activity proxy in the generated CSVs and figure hover data. That keeps the tables compact while preserving one explicit cost surface beyond route presence alone.

Routing-Fitness Crossover Summary

Compact crossover view for the remaining route-visible design question.

Question	Band	Engine Set	Route	Recov.	Loss	Churn	Hop
Maintenance benefit crossover	Low	BATMAN Classic	833	0	-	0.0	5.0
Maintenance benefit crossover	Low	BATMAN Bellman	933	0	-	0.0	5.0
Maintenance benefit crossover	Low	Babel	833	0	-	0.0	6.0
Maintenance benefit crossover	Low	OLSRv2	833	0	-	0.0	5.0
Maintenance benefit crossover	Low	Pathway	933	0	-	0.0	5.0
Maintenance benefit crossover	Low	Scatter	933	0	-	0.0	2.0
Maintenance benefit crossover	Low	Pathway + BATMAN Bellman	933	0	-	0.0	5.0
Maintenance benefit crossover	Low	Field	933	0	-	0.0	3.0
Maintenance benefit crossover	Moderate	BATMAN Classic	857	0	-	1.0	5.0
Maintenance benefit crossover	Moderate	BATMAN Bellman	928	0	-	1.0	5.0
Maintenance benefit crossover	Moderate	Babel	857	0	-	1.0	6.0
Maintenance benefit crossover	Moderate	OLSRv2	857	0	-	1.0	5.0
Maintenance benefit crossover	Moderate	Pathway	928	0	-	1.0	5.0
Maintenance benefit crossover	Moderate	Scatter	928	0	-	0.0	2.0
Maintenance benefit crossover	Moderate	Pathway + BATMAN Bellman	928	0	-	3.0	5.0

Question	Band	Engine Set	Route	Recov.	Loss	Churn	Hop
Maintenance benefit crossover	Moderate	Field	285	0	15	0.0	3.0
Maintenance benefit crossover	High	BATMAN Classic	840	0	-	3.0	6.0
Maintenance benefit crossover	High	BATMAN Bellman	920	0	-	1.0	6.0
Maintenance benefit crossover	High	Babel	840	0	-	1.0	10.0
Maintenance benefit crossover	High	OLSRv2	840	0	-	3.0	6.0
Maintenance benefit crossover	High	Pathway	920	0	-	1.0	7.0
Maintenance benefit crossover	High	Scatter	920	0	-	0.0	2.0
Maintenance benefit crossover	High	Pathway + BATMAN Bellman	920	0	-	7.0	6.0
Maintenance benefit crossover	High	Field	200	0	14	0.0	3.0
Search burden crossover	Low	BATMAN Classic	888	0	-	1.0	2.0
Search burden crossover	Low	BATMAN Bellman	888	0	-	0.0	2.0
Search burden crossover	Low	Babel	888	0	-	0.0	2.0
Search burden crossover	Low	OLSRv2	888	0	-	0.0	2.0
Search burden crossover	Low	Pathway	888	0	-	0.0	2.0
Search burden crossover	Low	Scatter	888	0	-	0.0	2.0
Search burden crossover	Low	Pathway + BATMAN Bellman	888	0	-	1.0	2.0
Search burden crossover	Low	Field	888	0	-	0.0	1.0
Search burden crossover	Moderate	BATMAN Classic	805	0	-	0.0	7.0
Search burden crossover	Moderate	BATMAN Bellman	916	0	-	0.0	7.0
Search burden crossover	Moderate	Babel	805	0	-	0.0	8.0
Search burden crossover	Moderate	OLSRv2	805	0	-	0.0	7.0
Search burden crossover	Moderate	Pathway	916	0	-	1.0	7.0
Search burden crossover	Moderate	Scatter	916	0	-	0.0	2.0
Search burden crossover	Moderate	Pathway + BATMAN Bellman	916	0	-	0.0	7.0

Question	Band	Engine Set	Route	Recov.	Loss	Churn	Hop
Search burden crossover	Moderate	Field	55	0	5	0.0	3.0
Search burden crossover	High	BATMAN Classic	772	0	-	0.0	10.0
Search burden crossover	High	BATMAN Bellman	909	0	-	0.0	10.0
Search burden crossover	High	Babel	772	0	-	0.0	12.0
Search burden crossover	High	OLSRv2	772	0	-	0.0	10.0
Search burden crossover	High	Pathway	909	0	-	1.0	10.0
Search burden crossover	High	Scatter	909	0	-	0.0	2.0
Search burden crossover	High	Pathway + BATMAN Bellman	909	0	-	0.0	10.0
Search burden crossover	High	Field	45	0	6	0.0	4.0

Table 4. Question is the analytical axis, Band is the maintained difficulty band, Engine Set is the standalone routing stack, Route is total-window route presence, Recov. is recovery success after a loss, Loss is the first-loss round, Churn is mean route churn, and Hop is the active-route hop-count proxy.

Routing-Fitness Multi-Flow Summary

Compact fairness view for the shared-broker families.

Family	Engine Set	Min	Max	Spread	Starved	Broker P/C/S	Live	Churn
Shared corridor	BATMAN Classic	928	1000	72	0	100/100/0.0	13.0	0.0
Shared corridor	BATMAN Bellman	1000	1000	0	0	100/100/0.0	14.0	0.0
Shared corridor	Babel	928	1000	72	0	100/100/0.0	13.0	0.0
Shared corridor	OLSRv2	928	1000	72	0	100/100/0.0	13.0	0.0
Shared corridor	Pathway	0	1000	1000	2	100/100/0.0	0.0	0.0
Shared corridor	Scatter	1000	1000	0	0	-	14.0	0.0
Shared corridor	Pathway + BATMAN Bellman	0	1000	1000	2	100/100/0.0	0.0	0.0
Shared corridor	Field	0	0	0	3	-	0.0	0.0
Asymmetric demand	BATMAN Classic	933	1000	67	0	100/100/0.0	28.0	0.0
Asymmetric demand	BATMAN Bellman	1000	1000	0	0	100/100/0.0	30.0	0.0
Asymmetric demand	Babel	933	1000	67	0	100/100/0.0	28.0	0.0

Family	Engine Set	Min	Max	Spread	Starved	Broker P/C/S	Live	Churn
Asymmetric demand	OLSRv2	933	1000	67	0	100/100/ 0.0	28.0	0.0
Asymmetric demand	Pathway	0	1000	1000	2	100/100/ 0.0	0.0	0.0
Asymmetric demand	Scatter	1000	1000	0	0	-	30.0	0.0
Asymmetric demand	Pathway + BATMAN Bellman	0	1000	1000	2	100/100/ 0.0	0.0	0.0
Asymmetric demand	Field	0	433	433	2	0/0/0.0	0.0	0.0
Detour choice	BATMAN Classic	964	1000	36	0	-	28.0	1.0
Detour choice	BATMAN Bellman	1000	1000	0	0	-	28.0	2.0
Detour choice	Babel	964	1000	36	0	-	27.0	3.0
Detour choice	OLSRv2	964	1000	36	0	-	27.0	0.0
Detour choice	Pathway	0	1000	1000	2	-	0.0	0.0
Detour choice	Scatter	1000	1000	0	0	-	28.0	0.0
Detour choice	Pathway + BATMAN Bellman	0	1000	1000	2	-	0.0	0.0
Detour choice	Field	0	0	0	3	-	0.0	0.0

Table 5. Min and Max are the weakest and strongest per-flow route-presence means, Spread is the gap between them, Starved is the mean count of objectives with zero route presence, Broker P/C/S reports broker participation percent, bottleneck concentration percent, and broker switch count as `participation/concentration/switches`, Live is the mean number of rounds where multiple objectives are simultaneously live, and Churn is mean route churn.

Routing-Fitness Stale Repair Summary

Compact stale-information repair view.

Family	Engine Set	Persist	Route	Unrec.	Loss	Churn
Delayed observation	BATMAN Classic	-	892	0	-	0.0
Delayed observation	BATMAN Bellman	-	928	0	-	0.0
Delayed observation	Babel	-	892	0	-	0.0
Delayed observation	OLSRv2	-	892	0	-	0.0
Delayed observation	Pathway	-	928	0	-	0.0
Delayed observation	Scatter	-	928	0	-	0.0
Delayed observation	Pathway + BATMAN Bellman	-	928	0	-	0.0
Delayed observation	Field	13	571	1	18	0.0
Asymmetric stale region	BATMAN Classic	-	892	0	-	0.0
Asymmetric stale region	BATMAN Bellman	-	928	0	-	0.0
Asymmetric stale region	Babel	-	892	0	-	0.0
Asymmetric stale region	OLSRv2	-	892	0	-	0.0
Asymmetric stale region	Pathway	-	928	0	-	0.0

Family	Engine Set	Persist	Route	Unrec.	Loss	Churn
Asymmetric stale region	Scatter	-	928	0	-	0.0
Asymmetric stale region	Pathway + BATMAN Bellman	-	928	0	-	0.0
Asymmetric stale region	Field	13	571	1	18	0.0
Recovery window	BATMAN Classic	-	900	0	-	0.0
Recovery window	BATMAN Bellman	-	933	0	-	0.0
Recovery window	Babel	-	900	0	-	0.0
Recovery window	OLSRv2	-	900	0	-	0.0
Recovery window	Pathway	-	933	0	-	0.0
Recovery window	Scatter	-	933	0	-	0.0
Recovery window	Pathway + BATMAN Bellman	-	933	0	-	0.0
Recovery window	Field	-	933	0	-	0.0

Table 6. Persist is the mean bad-route persistence after the first disruptive topology change, Route is total-window route presence, Unrec. is mean unrecovered-after-loss count, Loss is the first-loss round, and Churn is mean route churn. Recovery-event success is still exported in the CSV, but it is not used as the headline because many stale scenarios never enter a loss/recovery event path.

Routing-Fitness Figure Context

These figures isolate the last three decision questions directly: crossover under larger graph pressure, fairness under shared-broker contention, and stale-route persistence under delayed observations. They are meant to be read as envelope charts, not just winner charts.

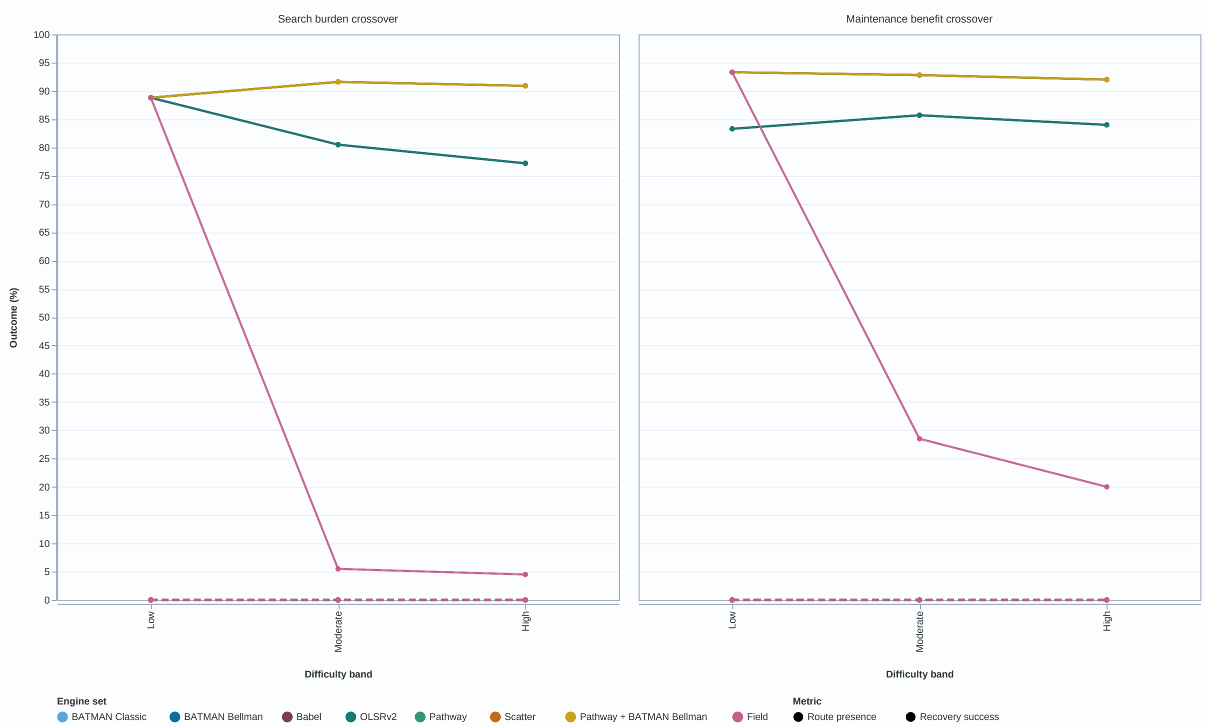

Figure 25. Crossover view for the remaining route-visible design boundary. Each panel fixes one analytical question and moves from low to high difficulty. Solid lines show total-window route presence and dashed lines show recovery success, so the useful operating envelope is where an engine keeps both high as the band hardens.

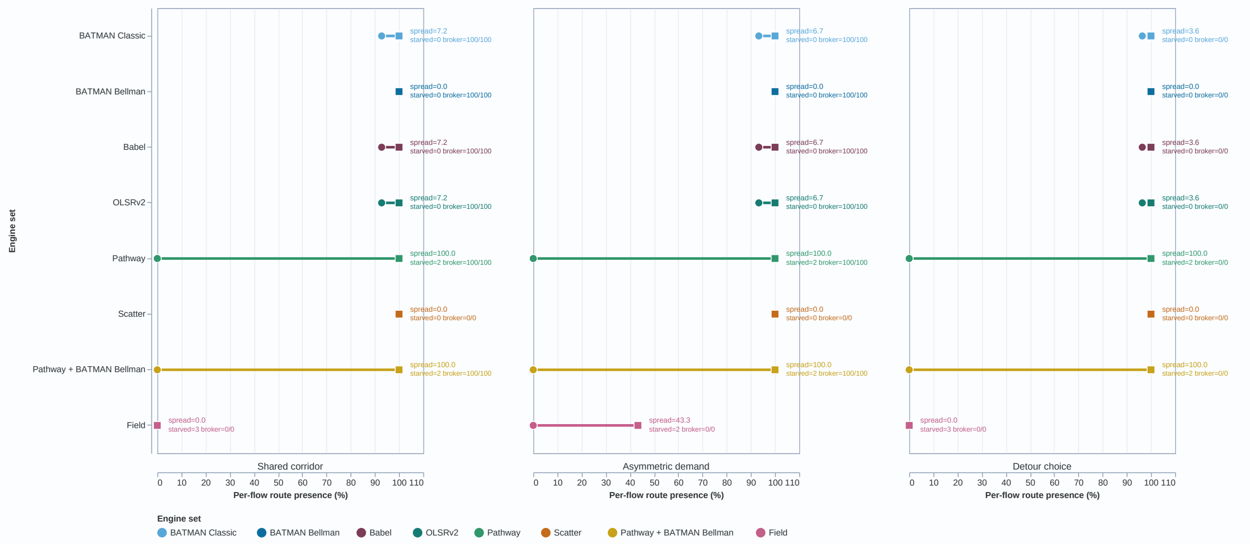

Figure 26. Multi-flow fairness under shared-broker contention. Each row spans the weakest-to-strongest per-flow route-presence results for one engine set in that family. Narrow spans with high left endpoints are good because they mean the weakest flow still gets service instead of the mean hiding a bad tail. The broker detail labels summarize how much of the visible route activity still traverses tagged brokers and how concentrated that load becomes on the hottest broker.

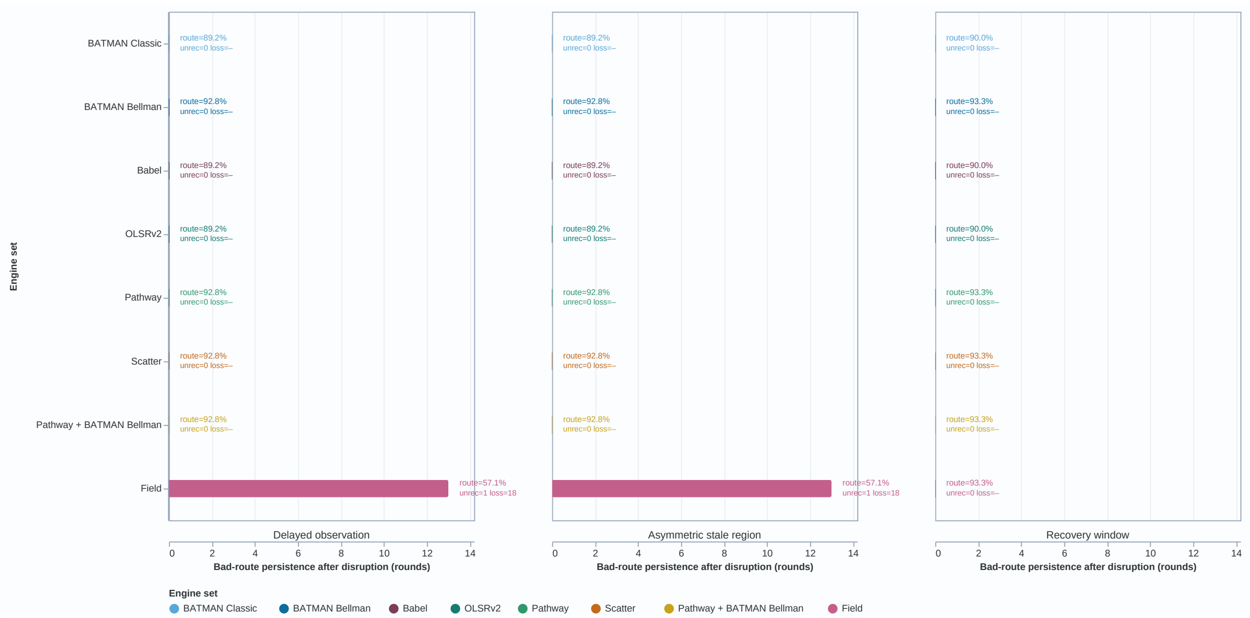

Figure 27. Bad-route persistence after delayed or asymmetric observations. Shorter bars are better because they mean the engine stops trusting stale routing state quickly after disruption. The labels show total-window route presence and unrecovered counts so the figure separates fast cleanup from cleanup that still leaves the objective unavailable.

Routing-Fitness Takeaways

- In the high search-burden crossover band, `batman-bellman`, `pathway`, `pathway-batman-bellman`, and `scatter` lead at 909 permille total-window route presence with 0 permille recovery success.
- In the high maintenance-benefit crossover band, `batman-bellman`, `pathway`, `pathway-batman-bellman`, and `scatter` lead at 920 permille total-window route presence with 0 permille recovery success.
- Under shared-broker contention, `Shared corridor` is best handled by `batman-bellman` and `scatter` at 1000 permille minimum per-flow route presence, while `Detour choice` is best handled

by `batman-bellman` and `scatter` at 1000 permille.

- The harshest fairness tail is currently `Shared corridor`, where `field` still records 3.0 starved objectives on average.
- In the stale-repair surface, `Recovery window` is best handled by `batman-bellman`, `field`, `pathway`, `pathway-batman-bellman`, and `scatter` at 0.0 rounds of stale persistence and 933 permille route presence, while the worst stale overcommit is `Delayed observation` under `babel` at 0.0 rounds and only 892 permille route presence.
- Taken together, the new evidence says the candidate direction is directionally supported, but still carrying one unresolved routing-risk regime.

Part III. Diffusion Calibration

13. Diffusion Calibration

Routing calibration and diffusion calibration are distinct objectives. Routing optimizes for activation, route presence, and recovery. Diffusion optimizes for eventual delivery, boundedness, latency, energy, and leakage.

For `field`, diffusion calibration also uses regime-specific success criteria: continuity families reward protected bridge-budget preservation and corridor persistence, scarcity families reward early conservative transition plus lower generic spread and expensive transport use, congestion families reward timely transition from cluster seeding into duplicate suppression without starving first-arrival cluster coverage, and privacy families reward lower observer leakage.

The maintained diffusion families are:

- `random-waypoint-sanity`: lightweight baseline with mixed movers.
- `partitioned-clusters`: separated clusters with rare bridger contacts.
- `disaster-broadcast`: urgent one-to-many over disrupted mobility.
- `sparse-long-delay`: sparse network with long delays and few long-range movers.
- `high-density-overload`: dense camp-like setting exposing overload.
- `mobility-shift`: clusters reconfiguring over time.
- `adversarial-observation`: clustered delivery with observer nodes exposing leakage.
- `bridge-drought`: prolonged low-contact bridge regime.
- `energy-starved-relay`: low-energy relay regime punishing over-forwarding.
- `congestion-cascade`: dense low-capacity broadcast regime.

Detailed diffusion calibration, `field`-calibration, and boundary tables are collected in Appendix C so the main comparison can stay focused on regime winners and the figure-level differences.

Part IV. Diffusion Engine Comparison

14. Diffusion Engine Comparison

This part evaluates the calibrated profiles directly. The emphasis shifts from admissibility to how the engines differ when diffusion itself is the comparison surface.

The comparison surface here is regime-aware rather than purely family-by-family. Continuity, scarcity, congestion, privacy, and balanced regimes reward different behaviors, so the first summary table reports the best engine set per regime before the full family matrix.

The generic family-by-family winner table is a representative weighted surface, not a universal truth. Appendix C includes both the maintained non-`field` baseline audit and a winner-sensitivity table showing where delivery-heavy and boundedness-heavy generic weights keep or change the family winner.

Diffusion Regime Comparison

Best-performing engine set per diffusion regime.

Regime	Engine Set	Delivery	Coverage	Cluster Cov.	Tx	State	Score
balanced	`field-privacy-search-4`	741.0	741.0	871.4	3.7	viable	1049.7
congestion	`transition-broad`	554.3	554.3	979.0	10.0	viable	1248.3
continuity	`batman-classic`	1000.0	1000.0	1000.0	3.0	viable	1553.8
privacy	`batman-classic`	1000.0	1000.0	1000.0	3.0	viable	1132.0
scarcity	`field-scarcity`	1000.0	1000.0	1000.0	2.0	viable	821.8

Table 7. Regime, Engine Set, Delivery, Coverage, Cluster Cov. (target-cluster coverage), Tx, State, Score.

Appendix C contains the full diffusion family matrix and the `field`-versus-best-alternative regime table.

Diffusion Figure Context

These two figures separate delivery success from resource boundedness across the most discriminating maintained diffusion families.

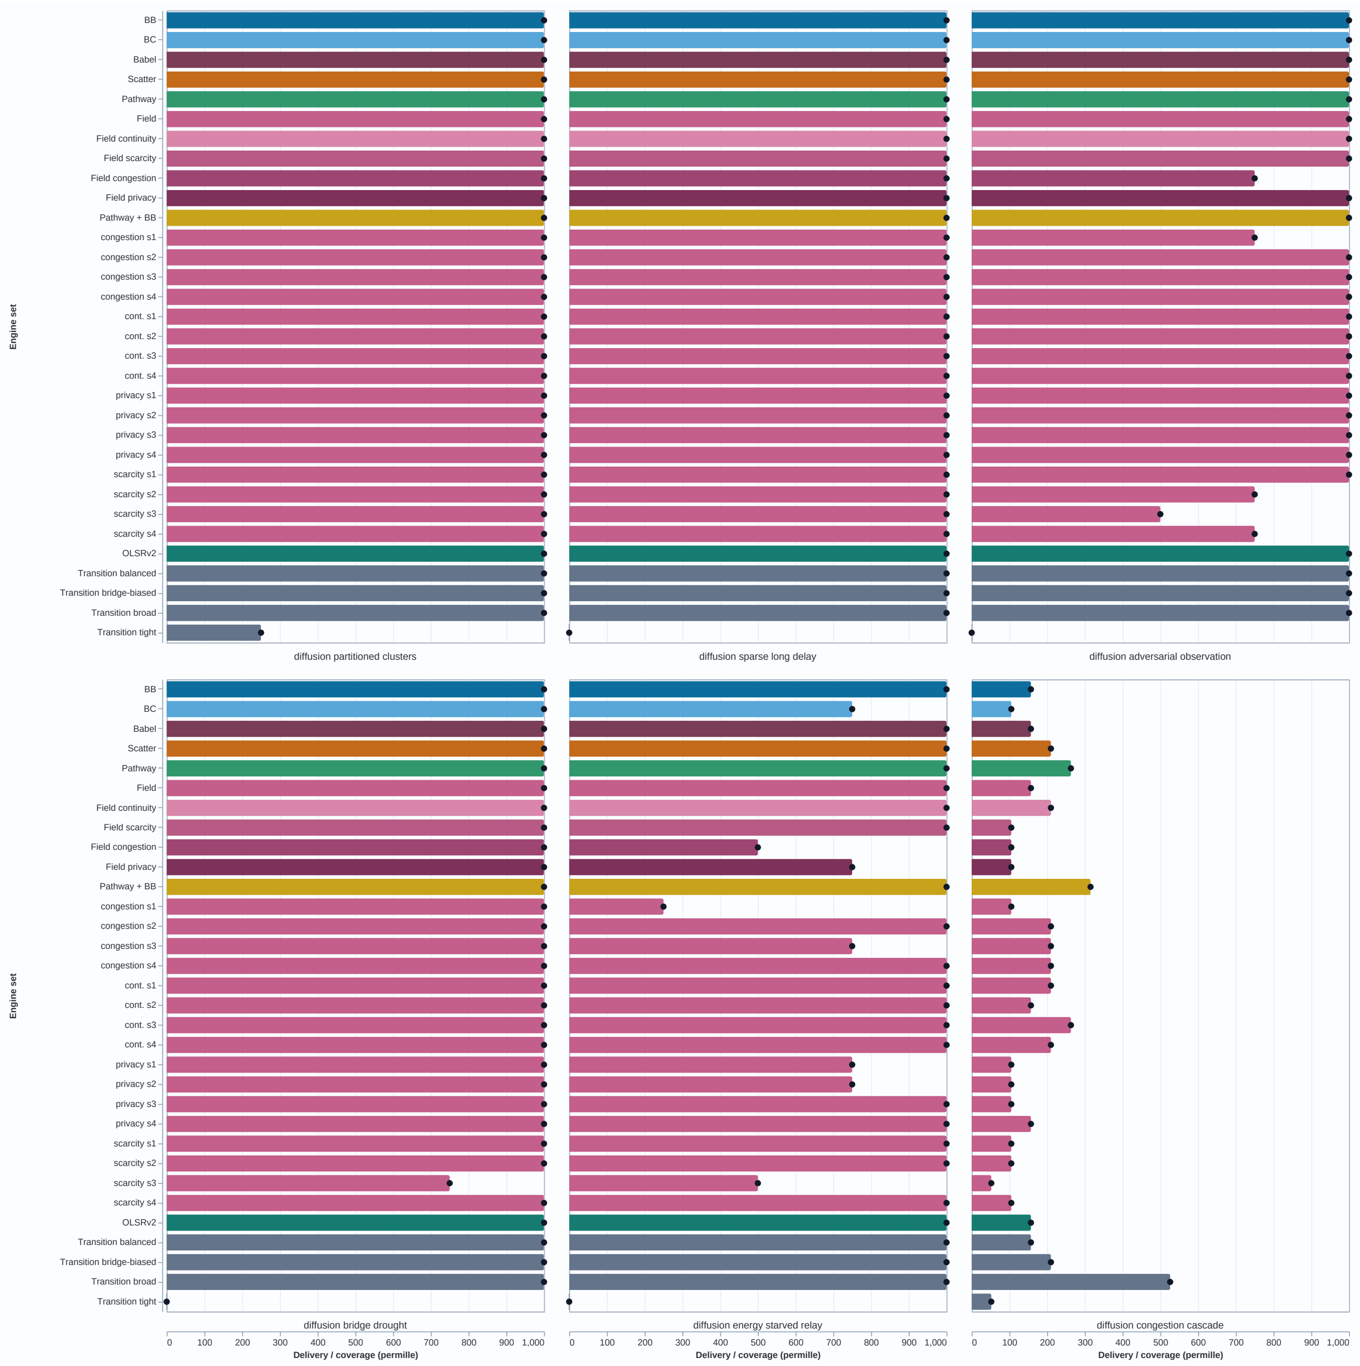

Figure 28. Diffusion delivery and coverage by scenario family. Longer bars are better because they indicate more successful delivery; the dot shows coverage, so a wider gap between bar and dot means delivery is lagging behind spread. Strong performers keep both high rather than trading one off against the other.

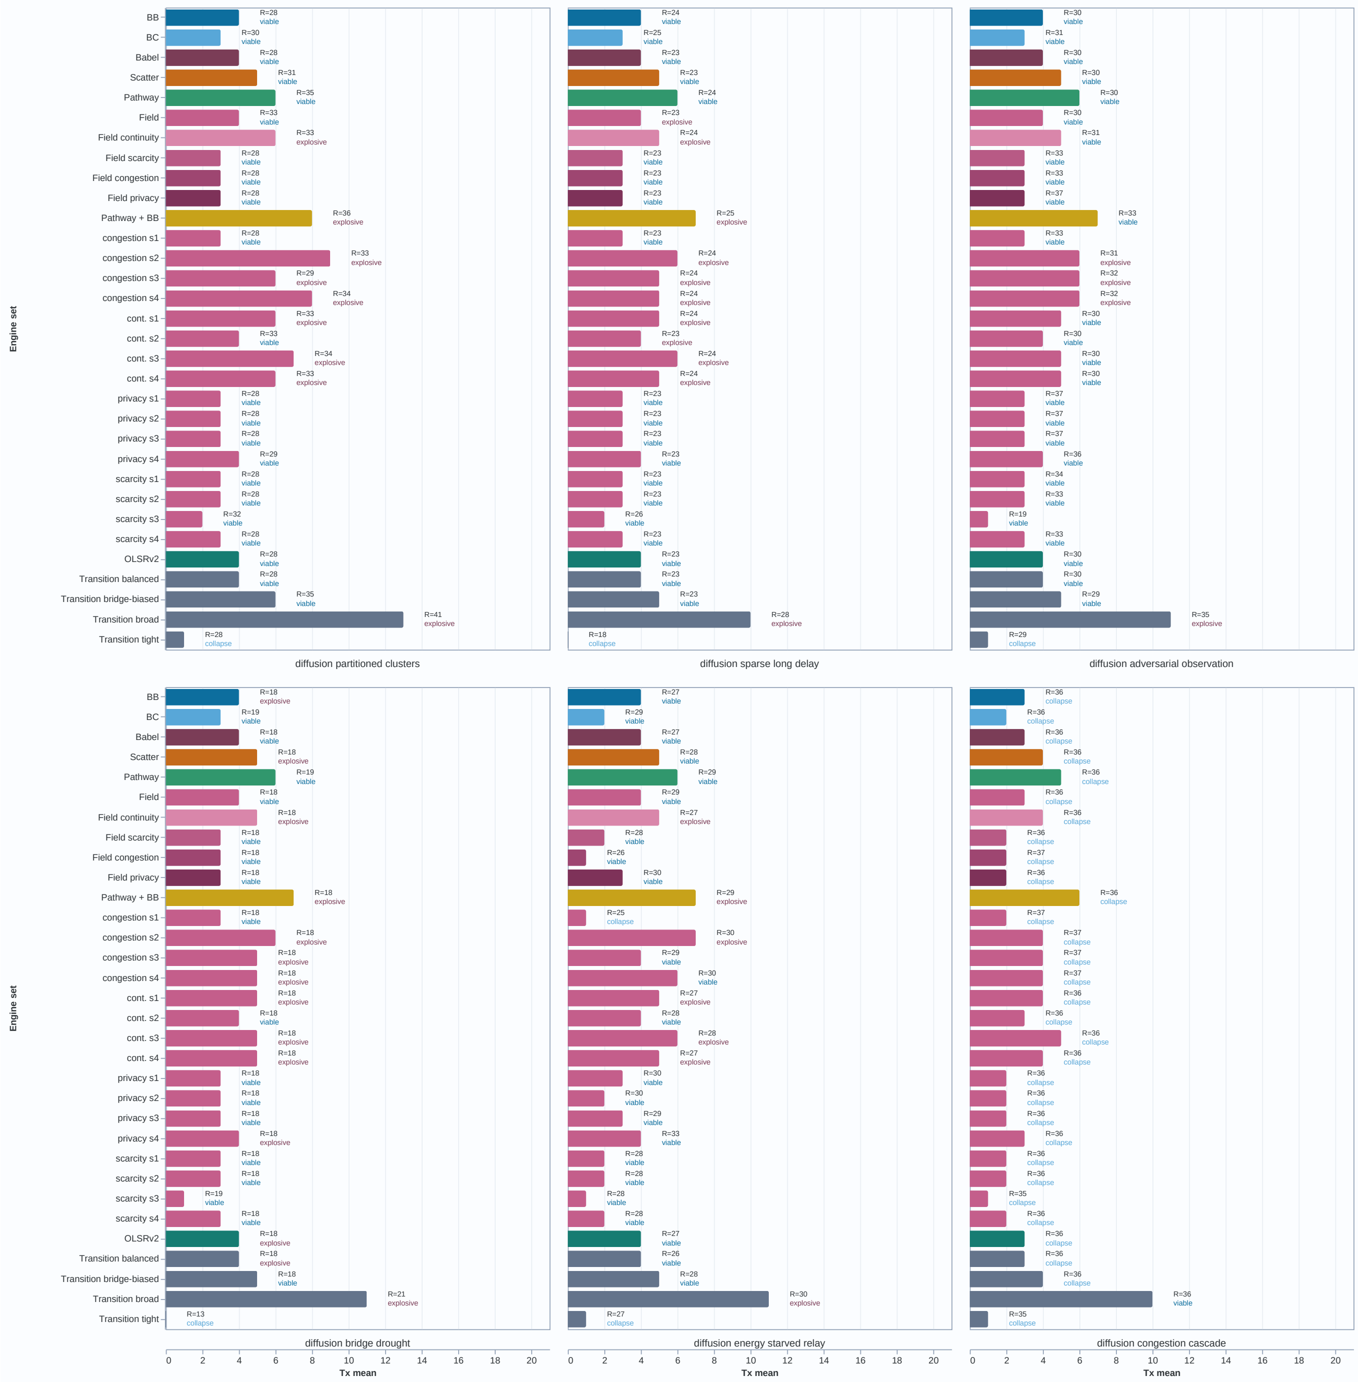

Figure 29. Diffusion transmission load and boundedness by scenario family. Lower transmission bars are better when delivery remains competitive because they indicate cheaper diffusion. The R= and bounded-state annotations show whether that load is staying inside the intended bounded operating regime or drifting toward over-spread.

Diffusion Takeaways

- The diffusion track is an engine comparison, but the regime summary is the right top-level view because continuity, scarcity, congestion, and privacy reward different tradeoffs.
- The regime winners currently split across concrete profiles rather than one family winner: `field-privacy-search-4` leads balanced, `field-scarcity` leads scarcity, `transition-broad` leads congestion, and `batman-classic` leads continuity and privacy.
- `field` shows clear regime specialization without being universal: it is `accepted` in balanced despite a regime-score delta of -472.1, beats the best alternative in scarcity by 70.1, matches privacy exactly at

0.0, remains acceptable but secondary in continuity at -67.3, and still has `no acceptable field candidate` in congestion.

- One under-represented point is that even the balanced regime currently prefers the privacy-tuned `field-privacy-search-4` profile, which suggests the maintained corpus is still rewarding bounded suppression-heavy behavior outside the explicit privacy regime.
- The harsher families are boundary finders: `bridge-drought` tests rare-opportunity carry, `energy-starved-relay` tests efficiency under scarcity, and `congestion-cascade` tests whether broad forwarding remains bounded without starving first-arrival cluster coverage.

The artifact set exposes posture-aware `field` diffusion behavior:

- In `bridge-drought`, `field` ends with dominant posture `continuity_biased` after 0 posture transitions, using 1 protected budget units and converting 1 of 18 protected bridge opportunities.
- In `energy-starved-relay`, `field` ends with dominant posture `scarcity_conservative`, first enters scarcity-conservative behavior at round 5, and suppresses 0 expensive transport attempts.
- In `congestion-cascade`, `field` ends with dominant posture `duplicate_suppressed`, first enters congestion control at round 3, seeds 3 first-arrival cluster transfers, records 260 cluster-coverage starvation events, and records 1540 redundant-forward suppressions plus 0 same-cluster suppressions.

Appendix A. Tuning Reference Tables

These tables provide the detailed tuning reference material behind the main recommendation and analysis sections.

Transition Behavior

This table shows how the leading configurations behave over time, not only how they score in aggregate.

Engine	Configuration	Route Mean	Route Stddev	First Mat.	First Loss	Recov.	Churn
`babel`	`babel-6-3`	922.3	18.8	2	-	-	0.0
`babel`	`babel-8-4`	922.3	18.8	2	-	-	0.0
`batman-bellman`	`batman-bellman-2-1`	780.0	142.5	2	8	7	0.9
`batman-bellman`	`batman-bellman-1-1`	780.0	142.5	2	8	7	0.7
`batman-classic`	`batman-classic-4-2`	653.0	97.3	12	17	10	0.7
`batman-classic`	`batman-classic-2-1`	639.3	97.4	12	17	11	0.7
comparison	`comparison-b6-3-p4-hop-lower-bound`	895.5	34.3	2	-	-	0.5
comparison	`comparison-b4-2-p3-zero`	517.7	196.5	2	-	-	0.6
`field`	`field-4-zero-p1-f140-n180`	780.9	199.8	3	14	-	0.2
`field`	`field-8-hop-lower-bound-p1-f120-n190`	780.9	199.8	3	14	-	0.2
`olsrv2`	`olsrv2-8-4`	792.2	134.7	2	8	7	1.0

Engine	Configuration	Route Mean	Route Stddev	First Mat.	First Loss	Recov.	Churn
olsrv2	olsrv2-6-3	781.8	149.1	2	8	7	1.2
pathway	pathway-2-hop-lower-bound	876.3	40.7	2	-	-	0.0
pathway	pathway-2-zero	876.3	40.7	2	-	-	0.0
scatter	scatter-conservative	900.5	22.4	2	-	-	0.0
scatter	scatter-balanced	900.5	22.4	2	-	-	0.0

Table 8. Route Mean and Route Stddev are average and spread of route presence across runs; First Mat. is the first round a route appears; First Loss is the first round a route disappears; Recovery is the first round routing returns after a loss; Churn counts route changes or handoffs. A - means the event was not observed.

Failure Boundaries

This table shows how much sustained stress each leading configuration survives before the first maintained failure.

Engine	Configuration	Max Stress	First Failed Family	Fail Stress	Reason
babel	babel-6-3	52	not observed	not observed	not observed
babel	babel-8-4	52	not observed	not observed	not observed
batman-bellman	batman-bellman-1-1	56	not observed	not observed	not observed
batman-bellman	batman-bellman-2-1	56	not observed	not observed	not observed
batman-classic	batman-classic-2-1	48	not observed	not observed	not observed
batman-classic	batman-classic-4-2	48	not observed	not observed	not observed
comparison	comparison-b6-3-p4-hop-lower-bound	90	not observed	not observed	not observed
comparison	comparison-b4-2-p3-zero	42	comparison-pathway-budget-boundary	44	activation-success
field	field-4-zero-p1-f140-n180	60	field-bridge-anti-entropy-continuity	52	route-presence
field	field-8-hop-lower-bound-p1-f120-n190	60	field-bridge-anti-entropy-continuity	52	route-presence
olsrv2	olsrv2-6-3	52	not observed	not observed	not observed
olsrv2	olsrv2-8-4	52	not observed	not observed	not observed
pathway	pathway-2-hop-lower-bound	58	not observed	not observed	not observed
pathway	pathway-2-zero	58	not observed	not observed	not observed
scatter	scatter-balanced	90	not observed	not observed	not observed

Engine	Configuration	Max Stress	First Failed Family	Fail Stress	Reason
scatter	scatter-conservative	90	not observed	not observed	not observed

Table 9. Max Stress is the highest stress level cleared; First Failed Family is the regime that breaks it; Fail Stress is the stress level at that failure; Reason is the dominant failure mode. `not observed` means no failure was observed.

Profile Recommendations

Engine	Profile	Configuration	Score	Activation	Route	Max Stress
batman-bellman	conservative	batman-bellman-1-1	5776.0	1000.0	780.0	56
batman-bellman	aggressive	batman-bellman-1-1	4025.9	1000.0	780.0	56
batman-bellman	degraded-network	batman-bellman-1-1	5812.7	1000.0	754.8	56
batman-classic	conservative	batman-classic-4-2	6675.2	1000.0	653.0	48
batman-classic	aggressive	batman-classic-4-2	4223.5	1000.0	653.0	48
batman-classic	degraded-network	batman-classic-4-2	6764.1	1000.0	713.0	48
babel	conservative	babel-8-4	6098.5	1000.0	922.3	52
babel	aggressive	babel-8-4	4365.5	1000.0	922.3	52
babel	degraded-network	babel-8-4	6033.9	1000.0	911.5	52
olsrv2	conservative	olsrv2-8-4	6717.9	1000.0	792.2	52
olsrv2	aggressive	olsrv2-8-4	4501.8	1000.0	792.2	52
olsrv2	degraded-network	olsrv2-8-4	6818.7	1000.0	821.3	52
scatter	balanced	scatter-balanced	5292.4	1000.0	900.5	90
scatter	conservative	scatter-conservative	7060.8	1000.0	900.5	90
scatter	degraded-network	scatter-degraded-network	7355.6	1000.0	901.5	90
pathway	balanced	pathway-2-hop-lower-bound	4166.3	1000.0	876.3	58
pathway	service-heavy	pathway-2-hop-lower-bound	5030.6	1000.0	876.3	58
pathway	degraded-network	pathway-2-hop-lower-bound	5771.4	1000.0	871.6	58
field	balanced	field-4-zero-p1-f140-n180	4606.5	1000.0	780.9	60
field	field-stable-service	field-4-zero-p1-f140-n180	4241.6	1000.0	780.9	60
field	field-low-churn	field-4-zero-p1-f140-n180	3453.7	1000.0	780.9	60

Engine	Profile	Configuration	Score	Activation	Route	Max Stress
`field`	`field`-broad-reselection	`field-4-zero-p3-f80-n70`	3881.2	1000.0	780.9	60
`field`	`field`-conservative-publication	`field-4-zero-p1-f140-n180`	3741.0	1000.0	780.9	60

Table 10. Profile is the ranking policy; Score is the profile-weighted composite value; Activation is share of runs that installed a route; Route is average route presence; Max Stress is highest stress level survived.

Field Continuity Profiles

This table treats Field lifecycle behavior as a tuning output in its own right.

Profile	Configuration	Score	Route	Shifts	Carry	Narrow	Degraded
`field`-stable-service	`field-4-zero-p1-f140-n180`	4241.6	780.9	1.7	6555.6	0.0	69.4
`field`-low-churn	`field-4-zero-p1-f140-n180`	3453.7	780.9	1.7	6555.6	0.0	69.4
`field`-broad-reselection	`field-4-zero-p3-f80-n70`	3881.2	780.9	8.4	13333.3	0.0	69.4
`field`-conservative-publication	`field-4-zero-p1-f140-n180`	3741.0	780.9	1.7	6555.6	0.0	69.4

Table 11. Profile names the continuity objective; Score is the profile-weighted value; Route is average route presence; Shifts is mean continuation-shift count; Carry is mean service carry-forward volume; Narrow is mean corridor-narrow count; Degraded is mean degraded-steady occupancy.

`field-stable-service` favors limited disruption; `field-low-churn` pushes harder against unnecessary movement; `field-broad-reselection` preserves more alternate branches and accepts more shifts; `field-conservative-publication` favors earlier narrowing and less corridor breadth.

Field Regime Calibration

This table calibrates `field` against regime-specific success criteria rather than only one flat route-visible recommendation score.

Regime	Success Criteria	Configuration	Route	Transition	Shifts	Carry	Stress
bootstrap-upgrade	upgrade bootstrap cleanly and avoid withdrawal or degraded fallback	`field-4-zero-p1-f140-n180`	879.5	-200.0	3.0	0.0	50
continuity-transition	retain corridor continuity through recovery with bounded shift pressure	`field-4-zero-p1-f140-n180`	662.7	86.7	1.3	0.0	54
service-continuity	preserve service continuity while keeping continuation churn bounded	`field-4-zero-p1-f140-n180`	820.2	130.0	1.2	14750.0	60

Table 12. Regime names the `field`-specific operating regime; Success Criteria states what the calibration is trying to optimize in that regime; Configuration is the best-scoring `field` setting for that regime; Route is mean route presence; Transition is the transition-health score; Shifts is mean continuation-shift count; Carry is mean service carry-forward volume; Stress is the highest stress envelope represented by the regime rows.

Appendix B. Route-Visible Reference Tables

These tables collect the exhaustive mixed-engine, mixed-engine selected-round breakdown, head-to-head route-visible results, and the remaining routing-fitness summary tables referenced by the main comparison sections.

Mixed-Engine Regime Split

Selected-round leader per maintained comparison family.

The mixed comparison surface is a single-router arbitration benchmark, not an oracle ensemble. The router gathers candidates across engines, publishes one canonical route per objective, and only reselects when maintenance or expiry requires it. Figure 15 therefore answers “which engine does the router actually use?” rather than “which engine is intrinsically best?” A mixed stack can legitimately underperform the best standalone engine in a family if the first durable admissible route comes from a weaker constituent engine.

Family	Selected-Round Leader	Activation	Active Route	Stress
comparison / bridge-transition	`batman-classic`	1000	1000	42
comparison / concurrent-mixed	`pathway`	1000	1000	48
comparison-connected / high-loss	`pathway`	1000	1000	54
comparison-connected / low-loss	`scatter`	1000	1000	18
comparison-corridor / continuity-uncertainty	`pathway`	1000	1000	50
comparison-large / core-periphery-high	`pathway`	1000	1000	76
comparison-large / core-periphery-moderate	`pathway`	1000	1000	66
comparison-large / multi-bottleneck-high	`pathway`	1000	1000	90
comparison-large / multi-bottleneck-moderate	`pathway`	1000	1000	82
comparison-medium / bridge-repair	`pathway`	1000	1000	58
comparison-multi / flow-asymmetric-demand	tie	1000	1000	74
comparison-multi / flow-detour-choice	tie	1000	1000	72
comparison-multi / flow-shared-corridor	tie	1000	1000	68
comparison-partial / observability-bridge	`olsrv2`	1000	1000	46
comparison-`pathway` / budget-boundary	`pathway`	1000	1000	44
comparison-stale / asymmetric-region	`pathway`	1000	1000	70
comparison-stale / observation-delay	`pathway`	1000	1000	62
comparison-stale / recovery-window	`pathway`	1000	1000	66

Table 13. Selected-Round Leader means the engine selected for the most active-route rounds in one shared mixed stack, not necessarily the best standalone performer. Activation is objective activation success, Active Route is active-window route presence, and Stress is the scenario stress score.

Mixed-Engine Selected-Round Breakdown

Each row reports the best maintained mixed comparison config for that family. The per-engine columns show average selected-route rounds under one shared router policy, so this table explains why the mixed stack leader is not an oracle best-of-engines result.

Family	Leader	Active Route	Handoffs	Batman Classic	Batman Bellman	Babel	OLSRv2	Scatter	Pathway	Field
comparison / bridge-transition	`batman-bellman`	1000	1	4	18	0	0	0	0	0
comparison / concurrent-mixed	`pathway`	1000	1	0	0	0	0	18	11	0
comparison-connected / high-loss	`pathway`	1000	3	0	7	6	1	20	0	0
comparison-connected / low-loss	`scatter`	1000	0	0	0	0	0	0	16	0
comparison-corridor / continuity-uncertainty	`pathway`	1000	0	0	0	0	0	25	0	0
comparison-large / core-periphery-high	`pathway`	1000	0	0	0	0	0	41	0	0
comparison-large / core-periphery-moderate	`pathway`	1000	0	0	0	0	0	33	0	0
comparison-large / multi-bottleneck-high	`pathway`	1000	0	0	0	0	0	47	0	0
comparison-large / multi-bottleneck-moderate	`pathway`	1000	0	0	0	0	0	39	0	0
comparison-medium / bridge-repair	`pathway`	1000	0	0	0	0	0	28	0	0
comparison-multi / flow-asymmetric-demand	tie	1000	0	0	0	0	0	30	30	0
comparison-multi / flow-detour-choice	tie	1000	0	0	0	0	0	28	28	0
comparison-multi / flow-shared-corridor	tie	1000	0	0	14	0	0	14	14	0
comparison-partial / observability-bridge	`olsrv2`	1000	0	0	0	0	21	20	0	0
comparison-`pathway` / budget-boundary	`pathway`	1000	0	0	0	0	0	8	0	0
comparison-stale / asymmetric-region	`pathway`	1000	0	0	0	0	0	26	0	0
comparison-stale / observation-delay	`pathway`	1000	0	0	0	0	0	26	0	0
comparison-stale / recovery-window	`pathway`	1000	0	0	0	0	0	28	0	0

Table 14. Family is the comparison regime. Leader is the selected-round leader. Active Route is active-window route presence. Handoffs is mean engine handoff count. The remaining columns show mean selected-route rounds per engine under the shared router policy.

Comparison Config Sensitivity Audit

Audit of whether the maintained configs separate each comparison family. `topline-and-selection` means both route outcomes and selected-engine behavior differ across configs. `selection-only` means route outcomes are identical but arbitration differs. `flat-control` means the family behaves identically under the maintained configs and should be read as a scenario/control row rather than a parameter-separation row.

Surface	Family	Class	Configs	Topline Sigs	Selection Sigs
comparison	comparison-connected / low-loss	flat-control	2	1	1
comparison	comparison / bridge-transition	topline-and-selection	2	2	2
comparison	comparison / concurrent-mixed	topline-and-selection	2	2	2
comparison	comparison-connected / high-loss	topline-and-selection	2	2	2
comparison	comparison-large / core-periphery-high	topline-and-selection	2	2	2
comparison	comparison-large / multi-bottleneck-high	topline-and-selection	2	2	2
comparison	comparison-multi / flow-shared-corridor	topline-and-selection	2	2	2
comparison	comparison-partial / observability-bridge	topline-and-selection	2	2	2
comparison	comparison-pathway / budget-boundary	topline-and-selection	2	2	2
comparison	comparison-corridor / continuity-uncertainty	topline-only	2	2	1
comparison	comparison-large / core-periphery-moderate	topline-only	2	2	1
comparison	comparison-large / multi-bottleneck-moderate	topline-only	2	2	1
comparison	comparison-medium / bridge-repair	topline-only	2	2	1
comparison	comparison-multi / flow-asymmetric-demand	topline-only	2	2	1
comparison	comparison-multi / flow-detour-choice	topline-only	2	2	1
comparison	comparison-stale / asymmetric-region	topline-only	2	2	1
comparison	comparison-stale / observation-delay	topline-only	2	2	1
comparison	comparison-stale / recovery-window	topline-only	2	2	1
head-to-head	head-to-head / partial-observability-bridge	selection-only	8	1	7
head-to-head	head-to / head-bridge-transition	topline-and-selection	8	5	7
head-to-head	head-to / head-concurrent-mixed	topline-and-selection	8	3	7
head-to-head	head-to-head / connected-high-loss	topline-and-selection	8	7	8
head-to-head	head-to-head / connected-low-loss	topline-and-selection	8	2	7

Surface	Family	Class	Configs	Topline Sigs	Selection Sigs
head-to-head	head-to-head / corridor-continuity-uncertainty	topline-and-selection	8	3	4
head-to-head	head-to-head / large-core-periphery-high	topline-and-selection	8	4	7
head-to-head	head-to-head / large-core-periphery-moderate	topline-and-selection	8	4	7
head-to-head	head-to-head / large-multi-bottleneck-high	topline-and-selection	8	6	7
head-to-head	head-to-head / large-multi-bottleneck-moderate	topline-and-selection	8	5	7
head-to-head	head-to-head / medium-bridge-repair	topline-and-selection	8	2	7
head-to-head	head-to-head / multi-flow-asymmetric-demand	topline-and-selection	8	4	7
head-to-head	head-to-head / multi-flow-detour-choice	topline-and-selection	8	7	7
head-to-head	head-to-head / multi-flow-shared-corridor	topline-and-selection	8	4	7
head-to-head	head-to-head / stale-asymmetric-region	topline-and-selection	8	3	7
head-to-head	head-to-head / stale-observation-delay	topline-and-selection	8	3	7
head-to-head	head-to-head / stale-recovery-window	topline-and-selection	8	2	7

Table 15. Surface is the comparison surface (`comparison` or `head-to-head`). Family is the regime. Class is the sensitivity classification. Configs is the number of configs observed. Topline Sigs and Selection Sigs are the counts of distinct outcome and arbitration signatures.

Benchmark Profile Audit

This appendix table separates the fixed representative benchmark configs used in the head-to-head surface from the calibrated-best profile recommendations in Part I. A `Match` row means the representative benchmark config and the current calibrated-best config happen to be the same.

Engine Set	Representative	Benchmark Config	Calibrated Profile	Calibrated Config	Match
`batman-classic`	fixed-representative	`head-to-head-batman-classic-4-2`	conservative	`batman-classic-4-2`	no
`batman-classic`	fixed-representative	`head-to-head-batman-classic-4-2`	aggressive	`batman-classic-4-2`	no
`batman-classic`	fixed-representative	`head-to-head-batman-classic-4-2`	degraded-network	`batman-classic-4-2`	no
`batman-bellman`	fixed-representative	`head-to-head-batman-bellman-1-1`	conservative	`batman-bellman-1-1`	no
`batman-bellman`	fixed-representative	`head-to-head-batman-bellman-1-1`	aggressive	`batman-bellman-1-1`	no
`batman-bellman`	fixed-representative	`head-to-head-batman-bellman-1-1`	degraded-network	`batman-bellman-1-1`	no

Engine Set	Representative	Benchmark Config	Calibrated Profile	Calibrated Config	Match
babel	fixed-representative	head-to-head-babel-4-2	conservative	babel-8-4	no
babel	fixed-representative	head-to-head-babel-4-2	aggressive	babel-8-4	no
babel	fixed-representative	head-to-head-babel-4-2	degraded-network	babel-8-4	no
olsrv2	fixed-representative	head-to-head-olsrv2-4-2	conservative	olsrv2-8-4	no
olsrv2	fixed-representative	head-to-head-olsrv2-4-2	aggressive	olsrv2-8-4	no
olsrv2	fixed-representative	head-to-head-olsrv2-4-2	degraded-network	olsrv2-8-4	no
pathway	fixed-representative	head-to-head-pathway-6-hop-lower-bound	balanced	pathway-2-hop-lower-bound	no
pathway	fixed-representative	head-to-head-pathway-6-hop-lower-bound	service-heavy	pathway-2-hop-lower-bound	no
pathway	fixed-representative	head-to-head-pathway-6-hop-lower-bound	degraded-network	pathway-2-hop-lower-bound	no
scatter	fixed-representative	head-to-head-scatter	balanced	scatter-balanced	no
scatter	fixed-representative	head-to-head-scatter	conservative	scatter-conservative	no
scatter	fixed-representative	head-to-head-scatter	degraded-network	scatter-degraded-network	no
pathway-batman-bellman	fixed-representative	head-to-head-pathway-batman-b6-3-p6-hop-lower-bound	-	-	no
field	fixed-representative	head-to-head-field-6-zero-p1-f140-n180	balanced	field-4-zero-p1-f140-n180	no
field	fixed-representative	head-to-head-field-6-zero-p1-f140-n180	field-stable-service	field-4-zero-p1-f140-n180	no
field	fixed-representative	head-to-head-field-6-zero-p1-f140-n180	field-low-churn	field-4-zero-p1-f140-n180	no
field	fixed-representative	head-to-head-field-6-zero-p1-f140-n180	field-broad-reselection	field-4-zero-p3-f80-n70	no
field	fixed-representative	head-to-head-field-6-zero-p1-f140-n180	field-conservative-publication	field-4-zero-p1-f140-n180	no

Table 16. Engine Set is the evaluated stack; Representative is the surface kind; Benchmark Config is the fixed benchmark configuration; Calibrated Profile is the calibrated-best profile; Calibrated Config is the calibrated-best configuration; Match indicates whether the benchmark and calibrated configurations are the same.

Head-To-Head Results

Direct stack-to-stack comparison: `batman-classic`, `batman-bellman`, `babel`, `olsrv2`, `scatter`, `pathway`, `pathway-batman-bellman`, and `field`.

Regime	Engine Set	Activation	Active Route	Selected Leader	Stress
head-to / head-bridge-transition	`pathway`	1000	1000	`pathway`	42
	`scatter`	1000	1000	`scatter`	42
	`field`	1000	1000	`field`	42
	`babel`	1000	954	`babel`	42
	`batman-classic`	1000	863	`batman-classic`	42
	`batman-bellman`	1000	818	`batman-bellman`	42
	`pathway-batman-bellman`	1000	818	`batman-bellman`	42
	`olsrv2`	1000	636	`olsrv2`	42
head-to / head-concurrent-mixed	`pathway`	1000	1000	`pathway`	48
	`scatter`	1000	1000	`scatter`	48
	`pathway-batman-bellman`	1000	1000	`pathway`	48
	`field`	1000	937	`field`	48
	`batman-classic`	500	500	`batman-classic`	48
	`batman-bellman`	500	500	`batman-bellman`	48
	`babel`	500	500	`babel`	48
	`olsrv2`	500	500	`olsrv2`	48
head-to-head / connected-high-loss	`pathway`	1000	1000	`pathway`	54
	`scatter`	1000	1000	`scatter`	54
	`pathway-batman-bellman`	1000	1000	`pathway`	54
	`field`	1000	909	`field`	54
	`babel`	1000	500	`babel`	54
	`batman-bellman`	1000	318	`batman-bellman`	54
	`olsrv2`	1000	318	`olsrv2`	54
	`batman-classic`	1000	272	`batman-classic`	54
head-to-head / connected-low-loss	`batman-classic`	1000	1000	`batman-classic`	18
	`batman-bellman`	1000	1000	`batman-bellman`	18
	`babel`	1000	1000	`babel`	18
	`olsrv2`	1000	1000	`olsrv2`	18

Regime	Engine Set	Activation	Active Route	Selected Leader	Stress
	pathway	1000	1000	pathway	18
	scatter	1000	1000	scatter	18
	pathway-batman-bellman	1000	1000	batman-bellman	18
	field	1000	1000	field	18
head-to-head / corridor-continuity-uncertainty	pathway	1000	1000	pathway	50
	scatter	1000	1000	scatter	50
	pathway-batman-bellman	1000	1000	pathway	50
	field	1000	320	field	50
	batman-classic	0	0	none	50
	batman-bellman	0	0	none	50
	babel	0	0	none	50
	olsrv2	0	0	none	50
head-to-head / large-core-periphery-high	batman-bellman	1000	1000	batman-bellman	76
	pathway	1000	1000	pathway	76
	scatter	1000	1000	scatter	76
	pathway-batman-bellman	1000	1000	batman-bellman	76
	batman-classic	1000	850	batman-classic	76
	babel	1000	850	babel	76
	olsrv2	1000	850	olsrv2	76
	field	1000	50	field	76
head-to-head / large-core-periphery-moderate	batman-bellman	1000	1000	batman-bellman	66
	pathway	1000	1000	pathway	66
	scatter	1000	1000	scatter	66
	pathway-batman-bellman	1000	1000	batman-bellman	66
	batman-classic	1000	878	batman-classic	66
	babel	1000	878	babel	66
	olsrv2	1000	878	olsrv2	66
	field	1000	60	field	66
head-to-head / large-multi-bottleneck-high	batman-bellman	1000	1000	batman-bellman	90
	pathway	1000	1000	pathway	90
	scatter	1000	1000	scatter	90

Regime	Engine Set	Activation	Active Route	Selected Leader	Stress
	pathway-batman-bellman	1000	1000	batman-bellman	90
	batman-classic	1000	913	batman-classic	90
	babel	1000	913	babel	90
	olsrv2	1000	913	olsrv2	90
	field	1000	217	field	90
head-to-head / large-multi-bottleneck-moderate	batman-bellman	1000	1000	batman-bellman	82
	pathway	1000	1000	pathway	82
	scatter	1000	1000	scatter	82
	pathway-batman-bellman	1000	1000	batman-bellman	82
	batman-classic	1000	923	batman-classic	82
	babel	1000	923	babel	82
	olsrv2	1000	923	olsrv2	82
	field	1000	307	field	82
head-to-head / medium-bridge-repair	batman-bellman	1000	1000	batman-bellman	58
	pathway	1000	1000	pathway	58
	scatter	1000	1000	scatter	58
	pathway-batman-bellman	1000	1000	batman-bellman	58
	field	1000	1000	field	58
	batman-classic	1000	892	batman-classic	58
	babel	1000	892	babel	58
	olsrv2	1000	892	olsrv2	58
head-to-head / multi-flow-asymmetric-demand	batman-bellman	1000	1000	batman-bellman	74
	scatter	1000	1000	scatter	74
	batman-classic	1000	966	batman-classic	74
	babel	1000	966	babel	74
	olsrv2	1000	966	olsrv2	74
	pathway	333	333	pathway	74
	pathway-batman-bellman	333	333	batman-bellman	74
	field	333	144	field	74
head-to-head / multi-flow-detour-choice	batman-bellman	1000	1000	batman-bellman	72

Regime	Engine Set	Activation	Active Route	Selected Leader	Stress
	`scatter`	1000	1000	`scatter`	72
	`batman-classic`	1000	988	`batman-classic`	72
	`babel`	1000	976	`babel`	72
	`olsrv2`	1000	976	`olsrv2`	72
	`pathway`	333	333	`pathway`	72
	`pathway-batman-bellman`	333	333	`batman-bellman`	72
	`field`	0	0	none	72
head-to-head / multi-flow-shared-corridor	`batman-bellman`	1000	1000	`batman-bellman`	68
	`scatter`	1000	1000	`scatter`	68
	`batman-classic`	1000	952	`batman-classic`	68
	`babel`	1000	952	`babel`	68
	`olsrv2`	1000	952	`olsrv2`	68
	`pathway`	333	333	`pathway`	68
	`pathway-batman-bellman`	333	333	`batman-bellman`	68
	`field`	0	0	none	68
head-to-head / partial-observability-bridge	`batman-classic`	1000	1000	`batman-classic`	46
	`batman-bellman`	1000	1000	`batman-bellman`	46
	`babel`	1000	1000	`babel`	46
	`olsrv2`	1000	1000	`olsrv2`	46
	`pathway`	1000	1000	`pathway`	46
	`scatter`	1000	1000	`scatter`	46
	`pathway-batman-bellman`	1000	1000	`batman-bellman`	46
	`field`	1000	1000	`field`	46
head-to-head / stale-asymmetric-region	`batman-bellman`	1000	1000	`batman-bellman`	70
	`pathway`	1000	1000	`pathway`	70
	`scatter`	1000	1000	`scatter`	70
	`pathway-batman-bellman`	1000	1000	`batman-bellman`	70
	`batman-classic`	1000	961	`batman-classic`	70
	`babel`	1000	961	`babel`	70
	`olsrv2`	1000	961	`olsrv2`	70
	`field`	1000	615	`field`	70

Regime	Engine Set	Activation	Active Route	Selected Leader	Stress
head-to-head / stale-observation-delay	`batman-bellman`	1000	1000	`batman-bellman`	62
	`pathway`	1000	1000	`pathway`	62
	`scatter`	1000	1000	`scatter`	62
	`pathway-batman-bellman`	1000	1000	`batman-bellman`	62
	`batman-classic`	1000	961	`batman-classic`	62
	`babel`	1000	961	`babel`	62
	`olsrv2`	1000	961	`olsrv2`	62
	`field`	1000	615	`field`	62
head-to-head / stale-recovery-window	`batman-bellman`	1000	1000	`batman-bellman`	66
	`pathway`	1000	1000	`pathway`	66
	`scatter`	1000	1000	`scatter`	66
	`pathway-batman-bellman`	1000	1000	`batman-bellman`	66
	`field`	1000	1000	`field`	66
	`batman-classic`	1000	964	`batman-classic`	66
	`babel`	1000	964	`babel`	66
	`olsrv2`	1000	964	`olsrv2`	66

Table 17. Engine Set is the only stack enabled; Activation is objective activation success; Active Route is active-window route presence; Selected Leader is the selected-round leader for that run and may be `tie`; Stress is the regime stress score.

Routing-Fitness Crossover Summary

Compact crossover view for the remaining route-visible design question.

Question	Band	Engine Set	Route	Recov.	Loss	Churn	Hop
Maintenance benefit crossover	Low	BATMAN Classic	833	0	-	0.0	5.0
Maintenance benefit crossover	Low	BATMAN Bellman	933	0	-	0.0	5.0
Maintenance benefit crossover	Low	Babel	833	0	-	0.0	6.0
Maintenance benefit crossover	Low	OLSRv2	833	0	-	0.0	5.0
Maintenance benefit crossover	Low	Pathway	933	0	-	0.0	5.0
Maintenance benefit crossover	Low	Scatter	933	0	-	0.0	2.0
Maintenance benefit crossover	Low	Pathway + BATMAN Bellman	933	0	-	0.0	5.0

Question	Band	Engine Set	Route	Recov.	Loss	Churn	Hop
Maintenance benefit crossover	Low	Field	933	0	-	0.0	3.0
Maintenance benefit crossover	Moderate	BATMAN Classic	857	0	-	1.0	5.0
Maintenance benefit crossover	Moderate	BATMAN Bellman	928	0	-	1.0	5.0
Maintenance benefit crossover	Moderate	Babel	857	0	-	1.0	6.0
Maintenance benefit crossover	Moderate	OLSRv2	857	0	-	1.0	5.0
Maintenance benefit crossover	Moderate	Pathway	928	0	-	1.0	5.0
Maintenance benefit crossover	Moderate	Scatter	928	0	-	0.0	2.0
Maintenance benefit crossover	Moderate	Pathway + BATMAN Bellman	928	0	-	3.0	5.0
Maintenance benefit crossover	Moderate	Field	285	0	15	0.0	3.0
Maintenance benefit crossover	High	BATMAN Classic	840	0	-	3.0	6.0
Maintenance benefit crossover	High	BATMAN Bellman	920	0	-	1.0	6.0
Maintenance benefit crossover	High	Babel	840	0	-	1.0	10.0
Maintenance benefit crossover	High	OLSRv2	840	0	-	3.0	6.0
Maintenance benefit crossover	High	Pathway	920	0	-	1.0	7.0
Maintenance benefit crossover	High	Scatter	920	0	-	0.0	2.0
Maintenance benefit crossover	High	Pathway + BATMAN Bellman	920	0	-	7.0	6.0
Maintenance benefit crossover	High	Field	200	0	14	0.0	3.0
Search burden crossover	Low	BATMAN Classic	888	0	-	1.0	2.0
Search burden crossover	Low	BATMAN Bellman	888	0	-	0.0	2.0
Search burden crossover	Low	Babel	888	0	-	0.0	2.0
Search burden crossover	Low	OLSRv2	888	0	-	0.0	2.0
Search burden crossover	Low	Pathway	888	0	-	0.0	2.0
Search burden crossover	Low	Scatter	888	0	-	0.0	2.0
Search burden crossover	Low	Pathway + BATMAN Bellman	888	0	-	1.0	2.0

Question	Band	Engine Set	Route	Recov.	Loss	Churn	Hop
Search burden crossover	Low	Field	888	0	-	0.0	1.0
Search burden crossover	Moderate	BATMAN Classic	805	0	-	0.0	7.0
Search burden crossover	Moderate	BATMAN Bellman	916	0	-	0.0	7.0
Search burden crossover	Moderate	Babel	805	0	-	0.0	8.0
Search burden crossover	Moderate	OLSRv2	805	0	-	0.0	7.0
Search burden crossover	Moderate	Pathway	916	0	-	1.0	7.0
Search burden crossover	Moderate	Scatter	916	0	-	0.0	2.0
Search burden crossover	Moderate	Pathway + BATMAN Bellman	916	0	-	0.0	7.0
Search burden crossover	Moderate	Field	55	0	5	0.0	3.0
Search burden crossover	High	BATMAN Classic	772	0	-	0.0	10.0
Search burden crossover	High	BATMAN Bellman	909	0	-	0.0	10.0
Search burden crossover	High	Babel	772	0	-	0.0	12.0
Search burden crossover	High	OLSRv2	772	0	-	0.0	10.0
Search burden crossover	High	Pathway	909	0	-	1.0	10.0
Search burden crossover	High	Scatter	909	0	-	0.0	2.0
Search burden crossover	High	Pathway + BATMAN Bellman	909	0	-	0.0	10.0
Search burden crossover	High	Field	45	0	6	0.0	4.0

Table 18. Question is the analytical axis, Band is the maintained difficulty band, Engine Set is the standalone routing stack, Route is total-window route presence, Recov. is recovery success after a loss, Loss is the first-loss round, Churn is mean route churn, and Hop is the active-route hop-count proxy.

Routing-Fitness Multi-Flow Summary

Compact fairness view for the shared-broker families.

Family	Engine Set	Min	Max	Spread	Starved	Broker P/C/S	Live	Churn
Shared corridor	BATMAN Classic	928	1000	72	0	100/100/0.0	13.0	0.0
Shared corridor	BATMAN Bellman	1000	1000	0	0	100/100/0.0	14.0	0.0

Family	Engine Set	Min	Max	Spread	Starved	Broker P/C/S	Live	Churn
Shared corridor	Babel	928	1000	72	0	100/100/0.0	13.0	0.0
Shared corridor	OLSRv2	928	1000	72	0	100/100/0.0	13.0	0.0
Shared corridor	Pathway	0	1000	1000	2	100/100/0.0	0.0	0.0
Shared corridor	Scatter	1000	1000	0	0	-	14.0	0.0
Shared corridor	Pathway + BATMAN Bellman	0	1000	1000	2	100/100/0.0	0.0	0.0
Shared corridor	Field	0	0	0	3	-	0.0	0.0
Asymmetric demand	BATMAN Classic	933	1000	67	0	100/100/0.0	28.0	0.0
Asymmetric demand	BATMAN Bellman	1000	1000	0	0	100/100/0.0	30.0	0.0
Asymmetric demand	Babel	933	1000	67	0	100/100/0.0	28.0	0.0
Asymmetric demand	OLSRv2	933	1000	67	0	100/100/0.0	28.0	0.0
Asymmetric demand	Pathway	0	1000	1000	2	100/100/0.0	0.0	0.0
Asymmetric demand	Scatter	1000	1000	0	0	-	30.0	0.0
Asymmetric demand	Pathway + BATMAN Bellman	0	1000	1000	2	100/100/0.0	0.0	0.0
Asymmetric demand	Field	0	433	433	2	0/0/0.0	0.0	0.0
Detour choice	BATMAN Classic	964	1000	36	0	-	28.0	1.0
Detour choice	BATMAN Bellman	1000	1000	0	0	-	28.0	2.0
Detour choice	Babel	964	1000	36	0	-	27.0	3.0
Detour choice	OLSRv2	964	1000	36	0	-	27.0	0.0
Detour choice	Pathway	0	1000	1000	2	-	0.0	0.0
Detour choice	Scatter	1000	1000	0	0	-	28.0	0.0
Detour choice	Pathway + BATMAN Bellman	0	1000	1000	2	-	0.0	0.0
Detour choice	Field	0	0	0	3	-	0.0	0.0

Table 19. Min and Max are the weakest and strongest per-flow route-presence means, Spread is the gap between them, Starved is the mean count of objectives with zero route presence, Broker P/C/S reports broker participation percent, bottleneck concentration percent, and broker switch count as `participation/concentration/switches`, Live is the mean number of rounds where multiple objectives are simultaneously live, and Churn is mean route churn.

Routing-Fitness Stale Repair Summary

Compact stale-information repair view.

Family	Engine Set	Persist	Route	Unrec.	Loss	Churn
Delayed observation	BATMAN Classic	-	892	0	-	0.0

Family	Engine Set	Persist	Route	Unrec.	Loss	Churn
Delayed observation	BATMAN Bellman	-	928	0	-	0.0
Delayed observation	Babel	-	892	0	-	0.0
Delayed observation	OLSRv2	-	892	0	-	0.0
Delayed observation	Pathway	-	928	0	-	0.0
Delayed observation	Scatter	-	928	0	-	0.0
Delayed observation	Pathway + BATMAN Bellman	-	928	0	-	0.0
Delayed observation	Field	13	571	1	18	0.0
Asymmetric stale region	BATMAN Classic	-	892	0	-	0.0
Asymmetric stale region	BATMAN Bellman	-	928	0	-	0.0
Asymmetric stale region	Babel	-	892	0	-	0.0
Asymmetric stale region	OLSRv2	-	892	0	-	0.0
Asymmetric stale region	Pathway	-	928	0	-	0.0
Asymmetric stale region	Scatter	-	928	0	-	0.0
Asymmetric stale region	Pathway + BATMAN Bellman	-	928	0	-	0.0
Asymmetric stale region	Field	13	571	1	18	0.0
Recovery window	BATMAN Classic	-	900	0	-	0.0
Recovery window	BATMAN Bellman	-	933	0	-	0.0
Recovery window	Babel	-	900	0	-	0.0
Recovery window	OLSRv2	-	900	0	-	0.0
Recovery window	Pathway	-	933	0	-	0.0
Recovery window	Scatter	-	933	0	-	0.0
Recovery window	Pathway + BATMAN Bellman	-	933	0	-	0.0
Recovery window	Field	-	933	0	-	0.0

Table 20. Persist is the mean bad-route persistence after the first disruptive topology change, Route is total-window route presence, Unrec. is mean unrecovered-after-loss count, Loss is the first-loss round, and Churn is mean route churn. Recovery-event success is still exported in the CSV, but it is not used as the headline because many stale scenarios never enter a loss/recovery event path.

Appendix C. Diffusion Reference Tables

These tables hold the exhaustive diffusion calibration and comparison material that supports the shorter regime-level discussion in the main body.

Diffusion Calibration Summary

Best-performing engine set per maintained diffusion family.

Here, `collapse` means the engine falls below the basic viability floor for delivery or coverage, while `explosive` means it preserves delivery only by driving reproduction, transmission, storage, or energy cost into an unbounded regime.

Family	Engine Set	Delivery	Coverage	Latency	State	Stress
diffusion / adversarial-observation	`field-scarcity`	1000	1000	5	viable	46
diffusion / bridge-drought	`field-congestion`	1000	1000	5	viable	72
diffusion / congestion-cascade	`transition-broad`	526	526	1	viable	78
diffusion / disaster-broadcast	`transition-broad`	549	549	1	viable	48
diffusion-energy / starved-relay	`field-scarcity-search-4`	1000	1000	6	viable	68
diffusion-high / density-overload	`transition-broad`	588	588	1	viable	62
diffusion-large / congestion-threshold-high	`transition-broad`	256	256	1	collapse	96
diffusion-large / congestion-threshold-moderate	`transition-broad`	370	370	1	viable	88
diffusion-large / regional-shift-high	`field-scarcity-search-3`	1000	1000	4	viable	84
diffusion-large / regional-shift-moderate	`field-privacy-search-2`	1000	1000	8	viable	76
diffusion-large / sparse-threshold-high	`field-scarcity-search-4`	1000	1000	5	viable	92
diffusion-large / sparse-threshold-moderate	`field-privacy-search-4`	1000	1000	9	viable	84
diffusion / mobility-shift	`field-scarcity-search-2`	1000	1000	6	viable	50
diffusion / partitioned-clusters	`field-scarcity-search-3`	1000	1000	7	viable	42
diffusion-random / waypoint-sanity	`field-scarcity-search-4`	1000	1000	5	viable	20
diffusion-sparse / long-delay	`field-scarcity-search-3`	1000	1000	6	viable	56

Table 21. Engine Set, Delivery (fraction of targets reached), Coverage (fraction of reachable nodes), Latency (mean delivery delay), State (boundedness classification), Stress.

Diffusion Baseline Audit

These rows summarize the maintained non-`field` diffusion baselines. They are representative benchmark configs, not a calibrated-best sweep, so the generic winner tables should be read with that scope in mind.

Config	Rep	TTL	Forward	Bridge	Delivery	Coverage	Cluster	State
`babel`	3	None	430	90	716.9	716.9	852.2	viable
`batman-bellman`	3	None	380	80	701.2	701.2	835.1	viable
`batman-classic`	2	None	320	60	639.4	639.4	739.2	viable
`olsrv2`	3	None	400	110	716.9	716.9	844.4	viable
`pathway`	5	None	540	180	747.1	747.1	914.2	viable

Config	Rep	TTL	Forward	Bridge	Delivery	Coverage	Cluster	State
`pathway-batman-bellman`	6	None	560	180	776.6	776.6	933.8	explosive
`scatter`	4	None	470	260	731.9	731.9	877.4	viable

Table 22. Config is the baseline configuration. Rep is the replication budget. TTL is the time-to-live in rounds. Forward is the forward probability. Bridge is the bridge bias. Delivery, Coverage, and Cluster are mean delivery, coverage, and cluster-coverage scores. State is the boundedness classification.

Diffusion Calibration Boundaries

Where each engine set stays viable and where it first collapses or becomes explosive.

Engine Set	Viable Families	First Collapse	Collapse Stress	First Explosive	Explosive Stress
`batman-classic`	11	diffusion-disaster-broadcast	48	-	-
`field-scarcity-search-3`	11	diffusion-disaster-broadcast	48	-	-
`babel`	10	diffusion-disaster-broadcast	48	diffusion-large-regional-shift-high	84
`field-congestion`	10	diffusion-disaster-broadcast	48	diffusion-large-regional-shift-high	84
`field-privacy`	10	diffusion-disaster-broadcast	48	diffusion-large-regional-shift-high	84
`field-privacy-search-1`	10	diffusion-disaster-broadcast	48	diffusion-large-regional-shift-high	84
`field-privacy-search-2`	10	diffusion-disaster-broadcast	48	diffusion-large-regional-shift-high	84
`field-privacy-search-3`	10	diffusion-disaster-broadcast	48	diffusion-large-regional-shift-high	84
`field-scarcity`	10	diffusion-disaster-broadcast	48	diffusion-large-regional-shift-high	84
`field-scarcity-search-1`	10	diffusion-disaster-broadcast	48	diffusion-large-regional-shift-high	84
`field-scarcity-search-2`	10	diffusion-disaster-broadcast	48	diffusion-large-regional-shift-high	84
`field-scarcity-search-4`	10	diffusion-disaster-broadcast	48	diffusion-large-regional-shift-high	84
`batman-bellman`	9	diffusion-disaster-broadcast	48	diffusion-bridge-drought	72
`field-congestion-search-1`	9	diffusion-disaster-broadcast	48	diffusion-large-regional-shift-high	84
`olsrv2`	9	diffusion-disaster-broadcast	48	diffusion-bridge-drought	72
`transition-balanced`	9	diffusion-disaster-broadcast	48	diffusion-bridge-drought	72
`transition-bridge-biased`	9	diffusion-disaster-broadcast	48	diffusion-large-regional-shift-high	84

Engine Set	Viable Families	First Collapse	Collapse Stress	First Explosive	Explosive Stress
`field-continuity-search-2`	8	diffusion-disaster-broadcast	48	diffusion-sparse-long-delay	56
`field-privacy-search-4`	8	diffusion-disaster-broadcast	48	diffusion-bridge-drought	72
`pathway`	8	diffusion-disaster-broadcast	48	diffusion-large-regional-shift-high	84
`scatter`	8	diffusion-disaster-broadcast	48	diffusion-bridge-drought	72
`field`	7	diffusion-disaster-broadcast	48	diffusion-sparse-long-delay	56
`pathway-batman-bellman`	6	diffusion-congestion-cascade	78	diffusion-partitioned-clusters	42
`field-congestion-search-3`	5	diffusion-high-density-overload	62	diffusion-partitioned-clusters	42
`field-continuity`	4	diffusion-disaster-broadcast	48	diffusion-partitioned-clusters	42
`field-continuity-search-1`	4	diffusion-disaster-broadcast	48	diffusion-partitioned-clusters	42
`field-continuity-search-4`	4	diffusion-disaster-broadcast	48	diffusion-partitioned-clusters	42
`transition-broad`	4	diffusion-large-congestion-threshold-high	96	diffusion-random-waypoint-sanity	20
`field-congestion-search-2`	3	diffusion-high-density-overload	62	diffusion-partitioned-clusters	42
`field-congestion-search-4`	3	diffusion-high-density-overload	62	diffusion-partitioned-clusters	42
`field-continuity-search-3`	3	diffusion-disaster-broadcast	48	diffusion-partitioned-clusters	42
`transition-tight`	0	diffusion-random-waypoint-sanity	20	-	-

Table 23. Viable Families, First Collapse, First Explosive.

Field Diffusion Regime Calibration

This table calibrates `field` on its own diffusion success surface instead of only asking whether the generic cross-engine score liked it.

Regime	Success Criteria	Configuration	Posture	State	Transition	Delivery	Tx	Fit
balanced	stay balanced when no stronger regime dominates	`field`	balanced	explosive	0.0 shifts	705.3	3.6	791.4
congestion	enter congestion suppression early enough to bound redundant spread with deterministic suppression memory	no acceptable (`field-congestion-search-2`)	duplicate_suppressed	collapse	congestion@4	303.3	5.0	643.7

Regime	Success Criteria	Configuration	Posture	State	Transition	Delivery	Tx	Fit
continuity	preserve protected bridge budget long enough to exploit rare continuity opportunities without overspread	`field-continuity-search-2`	continuity_biased	viable	0.0 shifts	1000.0	4.0	1481.5
privacy	reduce observer-adjacent dissemination while preserving delivery	`field-privacy`	balanced	viable	1.0 shifts	1000.0	3.0	1096.0
scarcity	enter scarcity early and cut generic spread, expensive transport use, and energy before explosiveness	`field-scarcity`	scarcity_conservative	viable	scarcity @5	1000.0	2.0	811.8

Table 24. Regime names the diffusion posture regime; Success Criteria states what `field` is supposed to optimize there; Configuration is the accepted `field` diffusion profile for that regime, or an explicit no-acceptable-candidate marker if every `field` candidate still fails; Posture is the dominant `field` posture; State is the dominant boundedness class; Transition summarizes either the posture transition count or the first scarcity / congestion transition round; Delivery is mean delivery success; Tx is mean transmission count; Fit is the regime-specific `field` fitness score. The CSV export also includes protected-budget use, bridge-opportunity capture, cluster-seeding use, coverage-starvation counts, and deterministic suppression counts for the winning profile or best attempt.

Field Vs Best Alternative

Best `field` candidate per diffusion regime against the best non-`field` alternative under the same regime-aware comparison score.

Regime	Field	OK	State	Alternative	Alt State	dDel	dCov	dClus	dTx	dScore
balanced	`field`	yes	explosive	`transition-bridge-biased`	viable	-44.7	-44.7	-45.1	-1.1	-472.1
congestion	`field-congestion-search-2`	no	collapse	`transition-broad`	viable	-251.0	-251.0	21.0	-5.0	-617.6
continuity	`field-continuity-search-2`	yes	viable	`batman-classic`	viable	0.0	0.0	0.0	1.0	-67.3
privacy	`field-privacy`	yes	viable	`batman-classic`	viable	0.0	0.0	0.0	0.0	0.0
scarcity	`field-scarcity`	yes	viable	`transition-balanced`	viable	0.0	0.0	0.0	-2.0	70.1

Table 25. Field is the best `field` attempt, OK reports whether that attempt cleared the `field`-specific acceptability gate, State / Alt State are boundedness modes, `dDel` / `dCov` / `dClus` are `field`-minus-alternative delivery, coverage, and target-cluster coverage deltas, `dTx` is transmission delta, and `dScore` is regime-score delta. Negative `dTx` is good for `field`.

Diffusion Winner Sensitivity

This table re-scores the generic Part IV family winners under delivery-heavy and boundedness-heavy weights. A `no` in Stable means the family-level winner is sensitive to generic weighting and should be read as provisional relative to the regime-specific tables.

Family	Balanced	Delivery-Heavy	Boundedness-Heavy	Stable
diffusion / adversarial-observation	`field-scarcity`	`field-scarcity`	`field-scarcity`	yes
diffusion / bridge-drought	`field-congestion`	`field-congestion`	`field-congestion`	yes
diffusion / congestion-cascade	`transition-broad`	`transition-broad`	`transition-broad`	yes
diffusion / disaster-broadcast	`transition-broad`	`transition-broad`	`transition-broad`	yes
diffusion-energy / starved-relay	`field-scarcity`	`field-scarcity`	`field-scarcity`	yes
diffusion-high / density-overload	`transition-broad`	`transition-broad`	`transition-broad`	yes
diffusion-large / congestion-threshold-high	`transition-broad`	`transition-broad`	`transition-broad`	yes
diffusion-large / congestion-threshold-moderate	`transition-broad`	`transition-broad`	`transition-broad`	yes
diffusion-large / regional-shift-high	`field-scarcity-search-3`	`field-scarcity-search-3`	`field-scarcity-search-3`	yes
diffusion-large / regional-shift-moderate	`field-congestion`	`field-congestion`	`field-congestion`	yes
diffusion-large / sparse-threshold-high	`field-scarcity-search-4`	`field-scarcity-search-4`	`field-scarcity-search-4`	yes
diffusion-large / sparse-threshold-moderate	`field-privacy-search-4`	`field-privacy-search-4`	`field-privacy-search-4`	yes
diffusion / mobility-shift	`field-scarcity-search-2`	`field`	`field-scarcity-search-2`	no
diffusion / partitioned-clusters	`field-scarcity-search-3`	`field-scarcity-search-3`	`field-scarcity-search-3`	yes
diffusion-random / waypoint-sanity	`field-privacy`	`field-privacy`	`field-privacy`	yes
diffusion-sparse / long-delay	`field-scarcity-search-3`	`field-scarcity-search-3`	`field-scarcity-search-3`	yes

Table 26. Family is the diffusion scenario. Balanced, Delivery-Heavy, and Boundedness-Heavy show the winning configuration under each weighting. Stable indicates whether the winner is consistent across all three weightings.

Diffusion Engine Comparison

Full maintained diffusion engine surface.

Family	Engine Set	Delivery	Coverage	Tx	R_est	State
diffusion / adversarial-observation	`field-scarcity`	1000	1000	3	33	viable
	`batman-classic`	1000	1000	3	31	viable
	`field`	1000	1000	4	30	viable
	`field-continuity-search-2`	1000	1000	4	30	viable
	`transition-bridge-biased`	1000	1000	5	29	viable
	`batman-bellman`	1000	1000	4	30	viable

Family	Engine Set	Delivery	Coverage	Tx	R_est	State
	`babel`	1000	1000	4	30	viable
	`olsrv2`	1000	1000	4	30	viable
	`transition-balanced`	1000	1000	4	30	viable
	`field-scarcity-search-1`	1000	1000	3	34	viable
	`field-privacy`	1000	1000	3	37	viable
	`field-privacy-search-1`	1000	1000	3	37	viable
	`field-privacy-search-2`	1000	1000	3	37	viable
	`field-privacy-search-3`	1000	1000	3	37	viable
	`field-continuity-search-1`	1000	1000	5	30	viable
	`field-continuity`	1000	1000	5	31	viable
	`scatter`	1000	1000	5	30	viable
	`field-continuity-search-4`	1000	1000	5	30	viable
	`field-privacy-search-4`	1000	1000	4	36	viable
	`pathway`	1000	1000	6	30	viable
	`field-continuity-search-3`	1000	1000	5	30	viable
	`pathway-batman-bellman`	1000	1000	7	33	viable
	`field-congestion`	750	750	3	33	viable
	`field-congestion-search-1`	750	750	3	33	viable
	`field-scarcity-search-2`	750	750	3	33	viable
	`field-scarcity-search-4`	750	750	3	33	viable
	`field-congestion-search-4`	1000	1000	6	32	explosive
	`field-congestion-search-2`	1000	1000	6	31	explosive
	`field-congestion-search-3`	1000	1000	6	32	explosive
	`transition-broad`	1000	1000	11	35	explosive
	`field-scarcity-search-3`	500	500	1	19	viable
	`transition-tight`	0	0	1	29	collapse
diffusion / bridge-drought	`field-congestion`	1000	1000	3	18	viable
	`field-congestion-search-1`	1000	1000	3	18	viable
	`field-privacy`	1000	1000	3	18	viable
	`field-privacy-search-1`	1000	1000	3	18	viable
	`field-privacy-search-2`	1000	1000	3	18	viable
	`field-privacy-search-3`	1000	1000	3	18	viable
	`field-scarcity`	1000	1000	3	18	viable
	`field-scarcity-search-1`	1000	1000	3	18	viable
	`field-scarcity-search-2`	1000	1000	3	18	viable
	`field-scarcity-search-4`	1000	1000	3	18	viable
	`batman-classic`	1000	1000	3	19	viable

Family	Engine Set	Delivery	Coverage	Tx	R_est	State
	babel	1000	1000	4	18	viable
	field	1000	1000	4	18	viable
	field-continuity-search-2	1000	1000	4	18	viable
	transition-bridge-biased	1000	1000	5	18	viable
	pathway	1000	1000	6	19	viable
	batman-bellman	1000	1000	4	18	explosive
	olsrv2	1000	1000	4	18	explosive
	transition-balanced	1000	1000	4	18	explosive
	field-scarcity-search-3	750	750	1	19	viable
	field-privacy-search-4	1000	1000	4	18	explosive
	scatter	1000	1000	5	18	explosive
	field-continuity	1000	1000	5	18	explosive
	field-continuity-search-1	1000	1000	5	18	explosive
	field-continuity-search-4	1000	1000	5	18	explosive
	pathway-batman-bellman	1000	1000	7	18	explosive
	field-congestion-search-3	1000	1000	5	18	explosive
	field-congestion-search-4	1000	1000	5	18	explosive
	field-continuity-search-3	1000	1000	5	18	explosive
	field-congestion-search-2	1000	1000	6	18	explosive
	transition-broad	1000	1000	11	21	explosive
	transition-tight	0	0	0	13	collapse
diffusion / congestion-cascade	transition-broad	526	526	10	36	viable
	pathway-batman-bellman	315	315	6	36	collapse
	pathway	263	263	5	36	collapse
	field-continuity-search-3	263	263	5	36	collapse
	field-congestion-search-2	210	210	4	37	collapse
	field-congestion-search-3	210	210	4	37	collapse
	field-congestion-search-4	210	210	4	37	collapse
	field-continuity	210	210	4	36	collapse
	field-continuity-search-1	210	210	4	36	collapse
	field-continuity-search-4	210	210	4	36	collapse
	scatter	210	210	4	36	collapse
	transition-bridge-biased	210	210	4	36	collapse
	field-continuity-search-2	157	157	3	36	collapse
	field-privacy-search-4	157	157	3	36	collapse
	field	157	157	3	36	collapse
	babel	157	157	3	36	collapse

Family	Engine Set	Delivery	Coverage	Tx	R_est	State
	olsrv2	157	157	3	36	collapse
	batman-bellman	157	157	3	36	collapse
	transition-balanced	157	157	3	36	collapse
	field-privacy	105	105	2	36	collapse
	field-privacy-search-1	105	105	2	36	collapse
	field-privacy-search-2	105	105	2	36	collapse
	field-privacy-search-3	105	105	2	36	collapse
	field-scarcity	105	105	2	36	collapse
	field-scarcity-search-1	105	105	2	36	collapse
	field-scarcity-search-2	105	105	2	36	collapse
	field-scarcity-search-4	105	105	2	36	collapse
	field-congestion	105	105	2	37	collapse
	field-congestion-search-1	105	105	2	37	collapse
	batman-classic	105	105	2	36	collapse
	field-scarcity-search-3	52	52	1	35	collapse
	transition-tight	52	52	1	35	collapse
diffusion / disaster-broadcast	transition-broad	549	549	10	37	viable
	field-congestion-search-2	466	466	7	35	viable
	field-congestion-search-4	432	432	7	33	viable
	field-congestion-search-3	383	383	6	32	viable
	pathway-batman-bellman	383	383	6	31	viable
	field-continuity-search-3	316	316	5	32	collapse
	pathway	333	333	5	30	collapse
	field-continuity	266	266	4	30	collapse
	field-continuity-search-1	266	266	4	30	collapse
	field-continuity-search-4	266	266	4	30	collapse
	transition-bridge-biased	266	266	4	30	collapse
	scatter	266	266	4	30	collapse
	field	200	200	3	30	collapse
	field-continuity-search-2	200	200	3	30	collapse
	field-privacy-search-4	200	200	3	30	collapse
	babel	200	200	3	30	collapse
	batman-bellman	200	200	3	30	collapse
	olsrv2	200	200	3	30	collapse
	transition-balanced	200	200	3	30	collapse
	field-congestion	133	133	2	29	collapse
	field-congestion-search-1	133	133	2	29	collapse

Family	Engine Set	Delivery	Coverage	Tx	R_est	State
	`field-privacy`	133	133	2	29	collapse
	`field-privacy-search-1`	133	133	2	29	collapse
	`field-privacy-search-2`	133	133	2	29	collapse
	`field-privacy-search-3`	133	133	2	29	collapse
	`field-scarcity`	133	133	2	29	collapse
	`field-scarcity-search-1`	133	133	2	29	collapse
	`field-scarcity-search-2`	133	133	2	29	collapse
	`field-scarcity-search-4`	133	133	2	29	collapse
	`batman-classic`	133	133	2	29	collapse
	`field-scarcity-search-3`	66	66	1	29	collapse
	`transition-tight`	66	66	1	29	collapse
diffusion-energy / starved-relay	`field-scarcity`	1000	1000	2	28	viable
	`field-scarcity-search-1`	1000	1000	2	28	viable
	`field-scarcity-search-4`	1000	1000	2	28	viable
	`field-scarcity-search-2`	1000	1000	2	28	viable
	`transition-balanced`	1000	1000	4	26	viable
	`babel`	1000	1000	4	27	viable
	`batman-bellman`	1000	1000	4	27	viable
	`olsrv2`	1000	1000	4	27	viable
	`field-privacy-search-3`	1000	1000	3	29	viable
	`field-continuity-search-2`	1000	1000	4	28	viable
	`field`	1000	1000	4	29	viable
	`scatter`	1000	1000	5	28	viable
	`transition-bridge-biased`	1000	1000	5	28	viable
	`field-congestion-search-4`	1000	1000	6	30	viable
	`pathway`	1000	1000	6	29	viable
	`field-privacy-search-4`	1000	1000	4	33	viable
	`field-privacy-search-2`	750	750	2	30	viable
	`field-privacy`	750	750	3	30	viable
	`field-privacy-search-1`	750	750	3	30	viable
	`field-continuity`	1000	1000	5	27	explosive
	`field-continuity-search-1`	1000	1000	5	27	explosive
	`field-continuity-search-4`	1000	1000	5	27	explosive
	`field-continuity-search-3`	1000	1000	6	28	explosive
	`batman-classic`	750	750	2	29	viable
	`field-congestion-search-3`	750	750	4	29	viable
	`pathway-batman-bellman`	1000	1000	7	29	explosive

Family	Engine Set	Delivery	Coverage	Tx	R_est	State
	field-congestion-search-2	1000	1000	7	30	explosive
	transition-broad	1000	1000	11	30	explosive
	field-scarcity-search-3	500	500	1	28	viable
	field-congestion	500	500	1	26	viable
	field-congestion-search-1	250	250	1	25	collapse
	transition-tight	0	0	1	27	collapse
diffusion-high / density-overload	transition-broad	588	588	10	40	viable
	pathway-batman-bellman	352	352	6	39	viable
	field-continuity-search-3	294	294	5	39	collapse
	pathway	294	294	5	39	collapse
	field-continuity-search-4	235	235	4	39	collapse
	field-congestion-search-2	234	234	4	41	collapse
	field-congestion-search-4	220	220	4	41	collapse
	field-continuity	235	235	4	39	collapse
	field-continuity-search-1	235	235	4	39	collapse
	field-congestion-search-3	205	205	3	42	collapse
	field-privacy-search-4	176	176	3	39	collapse
	scatter	235	235	4	39	collapse
	transition-bridge-biased	235	235	4	39	collapse
	field	176	176	3	39	collapse
	field-continuity-search-2	176	176	3	39	collapse
	babel	176	176	3	39	collapse
	batman-bellman	176	176	3	39	collapse
	olsrv2	176	176	3	39	collapse
	transition-balanced	176	176	3	39	collapse
	field-privacy	117	117	2	39	collapse
	field-privacy-search-1	117	117	2	39	collapse
	field-privacy-search-2	117	117	2	39	collapse
	field-privacy-search-3	117	117	2	39	collapse
	field-scarcity	117	117	2	39	collapse
	field-scarcity-search-1	117	117	2	39	collapse
	field-scarcity-search-2	117	117	2	39	collapse
	field-scarcity-search-4	117	117	2	39	collapse
	batman-classic	117	117	2	39	collapse
	field-congestion	117	117	2	40	collapse
	field-congestion-search-1	117	117	2	40	collapse
	field-scarcity-search-3	58	58	1	38	collapse

Family	Engine Set	Delivery	Coverage	Tx	R_est	State
	`transition-tight`	58	58	1	38	collapse
diffusion-large / congestion-threshold-high	`transition-broad`	256	256	10	28	collapse
	`field-congestion-search-2`	205	205	8	28	collapse
	`field-congestion-search-4`	179	179	7	28	collapse
	`field-congestion-search-3`	153	153	6	28	collapse
	`pathway-batman-bellman`	153	153	6	28	collapse
	`field-continuity-search-3`	128	128	5	28	collapse
	`pathway`	128	128	5	28	collapse
	`field-continuity`	102	102	4	28	collapse
	`field-continuity-search-1`	102	102	4	28	collapse
	`field-continuity-search-4`	102	102	4	28	collapse
	`scatter`	102	102	4	28	collapse
	`transition-bridge-biased`	102	102	4	28	collapse
	`babel`	76	76	3	28	collapse
	`batman-bellman`	76	76	3	28	collapse
	`field`	76	76	3	28	collapse
	`field-continuity-search-2`	76	76	3	28	collapse
	`field-privacy-search-4`	76	76	3	28	collapse
	`olsrv2`	76	76	3	28	collapse
	`transition-balanced`	76	76	3	28	collapse
	`batman-classic`	51	51	2	28	collapse
	`field-congestion`	51	51	2	28	collapse
	`field-congestion-search-1`	51	51	2	28	collapse
	`field-privacy`	51	51	2	28	collapse
	`field-privacy-search-1`	51	51	2	28	collapse
	`field-privacy-search-2`	51	51	2	28	collapse
	`field-privacy-search-3`	51	51	2	28	collapse
	`field-scarcity`	51	51	2	28	collapse
	`field-scarcity-search-1`	51	51	2	28	collapse
	`field-scarcity-search-2`	51	51	2	28	collapse
	`field-scarcity-search-4`	51	51	2	28	collapse
	`field-scarcity-search-3`	25	25	1	27	collapse
	`transition-tight`	25	25	1	27	collapse
diffusion-large / congestion-threshold-moderate	`transition-broad`	370	370	10	32	viable
	`field-congestion-search-2`	286	286	8	33	collapse
	`field-congestion-search-4`	259	259	7	32	collapse

Family	Engine Set	Delivery	Coverage	Tx	R_est	State
	field-congestion-search-3	222	222	6	32	collapse
	field-continuity-search-3	185	185	5	32	collapse
	pathway-batman-bellman	222	222	6	32	collapse
	pathway	185	185	5	32	collapse
	field-continuity	148	148	4	32	collapse
	field-continuity-search-1	148	148	4	32	collapse
	field-continuity-search-4	148	148	4	32	collapse
	field	111	111	3	31	collapse
	field-continuity-search-2	111	111	3	31	collapse
	field-privacy-search-4	111	111	3	31	collapse
	scatter	148	148	4	32	collapse
	transition-bridge-biased	148	148	4	32	collapse
	batman-bellman	111	111	3	31	collapse
	babel	111	111	3	31	collapse
	olsrv2	111	111	3	31	collapse
	transition-balanced	111	111	3	31	collapse
	field-congestion	74	74	2	31	collapse
	field-congestion-search-1	74	74	2	31	collapse
	field-privacy	74	74	2	31	collapse
	field-privacy-search-1	74	74	2	31	collapse
	field-privacy-search-2	74	74	2	31	collapse
	field-privacy-search-3	74	74	2	31	collapse
	field-scarcity	74	74	2	31	collapse
	field-scarcity-search-1	74	74	2	31	collapse
	field-scarcity-search-2	74	74	2	31	collapse
	field-scarcity-search-4	74	74	2	31	collapse
	batman-classic	74	74	2	31	collapse
	field-scarcity-search-3	37	37	1	31	collapse
	transition-tight	37	37	1	31	collapse
diffusion-large / regional-shift-high	field-scarcity-search-3	1000	1000	2	13	viable
	batman-classic	1000	1000	3	15	viable
	field-privacy	1000	1000	3	14	explosive
	field-privacy-search-1	1000	1000	3	14	explosive
	field-privacy-search-3	1000	1000	3	14	explosive
	field-scarcity	1000	1000	3	14	explosive
	field-privacy-search-2	1000	1000	3	14	explosive
	field-scarcity-search-1	1000	1000	3	14	explosive

Family	Engine Set	Delivery	Coverage	Tx	R_est	State
	field-congestion	1000	1000	3	14	explosive
	field-congestion-search-1	1000	1000	3	14	explosive
	field-scarcity-search-2	1000	1000	3	14	explosive
	field-scarcity-search-4	1000	1000	3	14	explosive
	field	1000	1000	4	14	explosive
	field-continuity-search-2	1000	1000	4	14	explosive
	field-privacy-search-4	1000	1000	4	14	explosive
	babel	1000	1000	4	14	explosive
	transition-balanced	1000	1000	4	14	explosive
	olsrv2	1000	1000	4	14	explosive
	scatter	1000	1000	5	14	explosive
	transition-bridge-biased	1000	1000	5	14	explosive
	batman-bellman	1000	1000	4	14	explosive
	field-continuity	1000	1000	5	14	explosive
	field-continuity-search-1	1000	1000	5	14	explosive
	field-continuity-search-4	1000	1000	5	14	explosive
	pathway	1000	1000	6	14	explosive
	field-continuity-search-3	1000	1000	6	14	explosive
	pathway-batman-bellman	1000	1000	7	14	explosive
	field-congestion-search-3	1000	1000	7	14	explosive
	field-congestion-search-4	1000	1000	8	14	explosive
	field-congestion-search-2	1000	1000	9	14	explosive
	transition-broad	1000	1000	11	15	explosive
	transition-tight	0	0	1	13	collapse
diffusion-large / regional-shift-moderate	field-congestion	1000	1000	3	18	viable
	field-congestion-search-1	1000	1000	3	18	viable
	field-privacy	1000	1000	3	18	viable
	field-privacy-search-1	1000	1000	3	18	viable
	field-privacy-search-2	1000	1000	3	18	viable
	field-scarcity	1000	1000	3	18	viable
	batman-bellman	1000	1000	4	18	viable
	olsrv2	1000	1000	4	18	viable
	transition-balanced	1000	1000	4	18	viable
	field-continuity-search-2	1000	1000	4	19	viable
	babel	1000	1000	4	20	viable
	transition-bridge-biased	1000	1000	5	19	viable
	field-privacy-search-3	1000	1000	3	19	viable

Family	Engine Set	Delivery	Coverage	Tx	R_est	State
	field-scarcity-search-1	1000	1000	3	19	viable
	field-scarcity-search-2	1000	1000	3	19	viable
	field-scarcity-search-4	1000	1000	3	19	viable
	pathway	1000	1000	6	21	viable
	scatter	1000	1000	5	18	viable
	pathway-batman-bellman	1000	1000	7	20	viable
	field-congestion-search-3	1000	1000	6	20	viable
	field-privacy-search-4	1000	1000	4	18	explosive
	batman-classic	750	750	2	18	viable
	field-continuity-search-1	1000	1000	5	19	explosive
	field-continuity-search-4	1000	1000	5	19	explosive
	field	1000	1000	4	22	explosive
	field-continuity	1000	1000	5	19	explosive
	field-continuity-search-3	1000	1000	6	19	explosive
	field-congestion-search-4	1000	1000	7	21	explosive
	field-congestion-search-2	1000	1000	7	21	explosive
	transition-broad	1000	1000	11	21	explosive
	field-scarcity-search-3	500	500	1	14	viable
	transition-tight	0	0	1	17	collapse
diffusion-large / sparse-threshold-high	field-scarcity-search-4	1000	1000	3	11	viable
	field-scarcity-search-2	1000	1000	3	11	viable
	field-privacy-search-3	1000	1000	3	11	viable
	field-privacy	1000	1000	3	11	viable
	field-privacy-search-1	1000	1000	3	11	viable
	field-privacy-search-2	1000	1000	3	11	viable
	field-scarcity-search-1	1000	1000	3	11	viable
	field-scarcity	1000	1000	3	11	viable
	field-continuity-search-2	1000	1000	4	11	viable
	transition-balanced	1000	1000	4	11	viable
	field	1000	1000	4	11	viable
	field-privacy-search-4	1000	1000	4	11	viable
	scatter	1000	1000	5	11	viable
	transition-bridge-biased	1000	1000	5	11	viable
	olsrv2	1000	1000	4	11	viable
	babel	1000	1000	4	11	viable
	field-continuity	1000	1000	5	11	viable
	field-continuity-search-1	1000	1000	5	11	viable

Family	Engine Set	Delivery	Coverage	Tx	R_est	State
	field-continuity-search-4	1000	1000	5	11	viable
	field-scarcity-search-3	750	750	1	10	viable
	batman-classic	750	750	2	11	viable
	pathway	1000	1000	6	11	explosive
	field-congestion-search-3	1000	1000	5	11	explosive
	batman-bellman	750	750	3	11	viable
	field-continuity-search-3	1000	1000	6	11	explosive
	pathway-batman-bellman	1000	1000	7	11	explosive
	field-congestion-search-4	1000	1000	6	11	explosive
	field-congestion-search-2	1000	1000	6	11	explosive
	transition-broad	1000	1000	11	10	explosive
	field-congestion-search-1	500	500	2	11	viable
	field-congestion	500	500	2	11	viable
	transition-tight	250	250	1	10	collapse
diffusion-large / sparse-threshold-moderate	field-privacy-search-4	1000	1000	4	15	viable
	transition-bridge-biased	1000	1000	5	15	explosive
	field-continuity	1000	1000	5	14	explosive
	field-continuity-search-1	1000	1000	5	14	explosive
	field-continuity-search-4	1000	1000	5	14	explosive
	field-congestion	750	750	2	15	viable
	field-congestion-search-1	750	750	2	15	viable
	field-privacy	750	750	2	15	viable
	field-privacy-search-1	750	750	2	15	viable
	field-privacy-search-2	750	750	2	15	viable
	field-privacy-search-3	750	750	2	15	viable
	field-scarcity	750	750	2	15	viable
	field-scarcity-search-1	750	750	2	15	viable
	field-scarcity-search-2	750	750	2	15	viable
	field-scarcity-search-4	750	750	2	15	viable
	babel	750	750	3	15	viable
	batman-bellman	750	750	3	15	viable
	olsrv2	750	750	3	15	viable
	transition-balanced	750	750	3	15	viable
	field-continuity-search-3	1000	1000	5	15	explosive
	pathway-batman-bellman	1000	1000	7	16	explosive
	field-congestion-search-4	1000	1000	6	15	explosive
	field-continuity-search-2	750	750	3	14	explosive

Family	Engine Set	Delivery	Coverage	Tx	R_est	State
	`field-scarcity-search-3`	500	500	1	16	viable
	`transition-broad`	1000	1000	11	15	explosive
	`scatter`	750	750	4	16	explosive
	`field`	750	750	3	15	explosive
	`batman-classic`	500	500	2	14	viable
	`pathway`	750	750	5	15	explosive
	`field-congestion-search-2`	750	750	6	15	explosive
	`transition-tight`	0	0	1	14	collapse
	`field-congestion-search-3`	500	500	5	15	explosive
diffusion / mobility-shift	`field-scarcity-search-2`	1000	1000	3	26	viable
	`field-scarcity-search-4`	1000	1000	3	26	viable
	`field`	1000	1000	4	25	viable
	`field-continuity-search-2`	1000	1000	4	25	viable
	`field-privacy-search-4`	1000	1000	4	25	viable
	`field-congestion`	1000	1000	3	26	viable
	`field-congestion-search-1`	1000	1000	3	26	viable
	`field-privacy`	1000	1000	3	26	viable
	`field-privacy-search-1`	1000	1000	3	26	viable
	`field-privacy-search-2`	1000	1000	3	26	viable
	`field-privacy-search-3`	1000	1000	3	26	viable
	`field-scarcity`	1000	1000	3	26	viable
	`field-scarcity-search-1`	1000	1000	3	26	viable
	`batman-classic`	1000	1000	3	27	viable
	`transition-bridge-biased`	1000	1000	5	27	viable
	`babel`	1000	1000	4	27	viable
	`batman-bellman`	1000	1000	4	27	viable
	`olsrv2`	1000	1000	4	27	viable
	`transition-balanced`	1000	1000	4	27	viable
	`field-continuity`	1000	1000	5	26	viable
	`field-continuity-search-1`	1000	1000	5	26	viable
	`field-continuity-search-4`	1000	1000	5	26	viable
	`pathway`	1000	1000	6	25	viable
	`scatter`	1000	1000	5	26	viable
	`pathway-batman-bellman`	1000	1000	7	27	viable
	`field-continuity-search-3`	1000	1000	5	26	viable
	`field-congestion-search-2`	1000	1000	6	27	viable
	`field-scarcity-search-3`	750	750	1	27	viable

Family	Engine Set	Delivery	Coverage	Tx	R_est	State
	`field-congestion-search-4`	1000	1000	7	27	explosive
	`field-congestion-search-3`	750	750	6	28	viable
	`transition-broad`	1000	1000	11	31	explosive
	`transition-tight`	0	0	1	25	collapse
diffusion / partitioned-clusters	`field-scarcity-search-3`	1000	1000	2	32	viable
	`field-privacy`	1000	1000	3	28	viable
	`field-privacy-search-1`	1000	1000	3	28	viable
	`field-privacy-search-2`	1000	1000	3	28	viable
	`field-privacy-search-3`	1000	1000	3	28	viable
	`field-scarcity`	1000	1000	3	28	viable
	`field-scarcity-search-1`	1000	1000	3	28	viable
	`field-scarcity-search-2`	1000	1000	3	28	viable
	`field-scarcity-search-4`	1000	1000	3	28	viable
	`batman-classic`	1000	1000	3	30	viable
	`field-congestion`	1000	1000	3	28	viable
	`field-congestion-search-1`	1000	1000	3	28	viable
	`babel`	1000	1000	4	28	viable
	`batman-bellman`	1000	1000	4	28	viable
	`olsrv2`	1000	1000	4	28	viable
	`transition-balanced`	1000	1000	4	28	viable
	`field-privacy-search-4`	1000	1000	4	29	viable
	`scatter`	1000	1000	5	31	viable
	`field`	1000	1000	4	33	viable
	`field-continuity-search-2`	1000	1000	4	33	viable
	`transition-bridge-biased`	1000	1000	6	35	viable
	`pathway`	1000	1000	6	35	viable
	`field-continuity`	1000	1000	6	33	explosive
	`field-continuity-search-1`	1000	1000	6	33	explosive
	`field-continuity-search-4`	1000	1000	6	33	explosive
	`pathway-batman-bellman`	1000	1000	8	36	explosive
	`field-continuity-search-3`	1000	1000	7	34	explosive
	`field-congestion-search-3`	1000	1000	6	29	explosive
	`field-congestion-search-4`	1000	1000	8	34	explosive
	`field-congestion-search-2`	1000	1000	9	33	explosive
	`transition-broad`	1000	1000	13	41	explosive
	`transition-tight`	250	250	1	28	collapse

Family	Engine Set	Delivery	Coverage	Tx	R_est	State
diffusion-random / waypoint-sanity	`field-privacy`	1000	1000	3	50	viable
	`field-privacy-search-1`	1000	1000	3	50	viable
	`field-privacy-search-2`	1000	1000	3	50	viable
	`field-privacy-search-3`	1000	1000	3	50	viable
	`field-scarcity`	1000	1000	3	50	viable
	`field-scarcity-search-1`	1000	1000	3	50	viable
	`field-scarcity-search-2`	1000	1000	3	50	viable
	`field-scarcity-search-4`	1000	1000	3	50	viable
	`field-congestion`	1000	1000	3	51	viable
	`field-congestion-search-1`	1000	1000	3	51	viable
	`batman-classic`	1000	1000	3	48	viable
	`babel`	1000	1000	4	47	viable
	`field-scarcity-search-3`	1000	1000	2	74	viable
	`field`	1000	1000	4	49	viable
	`batman-bellman`	1000	1000	4	48	viable
	`olsrv2`	1000	1000	4	48	viable
	`transition-balanced`	1000	1000	4	48	viable
	`field-continuity-search-2`	1000	1000	4	49	viable
	`scatter`	1000	1000	5	48	viable
	`transition-bridge-biased`	1000	1000	5	50	viable
	`field-privacy-search-4`	1000	1000	4	50	viable
	`pathway`	1000	1000	6	49	viable
	`field-continuity`	1000	1000	5	45	viable
	`field-continuity-search-1`	1000	1000	5	45	viable
	`field-continuity-search-4`	1000	1000	5	45	viable
	`pathway-batman-bellman`	1000	1000	7	52	viable
	`field-continuity-search-3`	1000	1000	5	50	viable
	`field-congestion-search-4`	1000	1000	6	52	viable
	`field-congestion-search-3`	1000	1000	5	56	viable
	`field-congestion-search-2`	1000	1000	6	52	viable
	`transition-broad`	1000	1000	11	68	explosive
	`transition-tight`	0	0	1	41	collapse
diffusion-sparse / long-delay	`field-scarcity-search-3`	1000	1000	2	26	viable
	`field-congestion`	1000	1000	3	23	viable
	`field-congestion-search-1`	1000	1000	3	23	viable
	`field-privacy`	1000	1000	3	23	viable

Family	Engine Set	Delivery	Coverage	Tx	R_est	State
	`field-privacy-search-1`	1000	1000	3	23	viable
	`field-privacy-search-2`	1000	1000	3	23	viable
	`field-privacy-search-3`	1000	1000	3	23	viable
	`field-scarcity`	1000	1000	3	23	viable
	`field-scarcity-search-1`	1000	1000	3	23	viable
	`field-scarcity-search-2`	1000	1000	3	23	viable
	`field-scarcity-search-4`	1000	1000	3	23	viable
	`batman-classic`	1000	1000	3	25	viable
	`olsrv2`	1000	1000	4	23	viable
	`transition-balanced`	1000	1000	4	23	viable
	`babel`	1000	1000	4	23	viable
	`scatter`	1000	1000	5	23	viable
	`batman-bellman`	1000	1000	4	24	viable
	`transition-bridge-biased`	1000	1000	5	23	viable
	`field-privacy-search-4`	1000	1000	4	23	viable
	`pathway`	1000	1000	6	24	viable
	`field`	1000	1000	4	23	explosive
	`field-continuity-search-2`	1000	1000	4	23	explosive
	`pathway-batman-bellman`	1000	1000	7	25	explosive
	`field-continuity`	1000	1000	5	24	explosive
	`field-continuity-search-1`	1000	1000	5	24	explosive
	`field-continuity-search-4`	1000	1000	5	24	explosive
	`field-congestion-search-3`	1000	1000	5	24	explosive
	`field-congestion-search-4`	1000	1000	5	24	explosive
	`field-continuity-search-3`	1000	1000	6	24	explosive
	`field-congestion-search-2`	1000	1000	6	24	explosive
	`transition-broad`	1000	1000	10	28	explosive
	`transition-tight`	0	0	0	18	collapse

Table 27. Family, Engine Set, Delivery, Coverage, Tx (transmission count), `R_est` (boundedness signal), State (boundedness classification). Use this table to inspect family-level exceptions after the regime summary above.