A scheduling experiment with explicit constraints

Question

How can an operating application connect asset evidence, feasible maintenance and responsible authority? The project combined a synthetic scheduling investigation, a public fleet analysis, domain standards and a local authority-control exercise. These evidence populations answer different questions and remain separate.

Model element	Declared setting
Assets	Three service nodes; at least two healthy per modeled slot
Work	Two chains: drain / fan replacement / return and drain / patch / return
Resources	Two technicians with different skills, one scarce kit and timed fan receipts
Time	Six-hour horizon in half-hour integer slots
Downtime	Continuous from drain start through return, including waiting
Unknown critical evidence	Hold rather than invented data or a claim of impossibility

Competent comparators

The direct planner constructed earliest-feasible chains in both orders. Richer workflows used ordinary mixed-integer optimization with separate checking. An exhaustive enumerator supplied a reference for the restricted fixtures. Its implementation was separate, but the same author selected the assumptions and wrote the implementations. [P1]

Research design repair

Review found that the original 24-row comparison contained only sixteen distinct modeled states and overlapped development states. The original record was retained, the exposure claim withdrawn, and a new confirmation set excluded previously used states. This is a concrete correction to the evidentiary design.

What the scheduling comparison actually found

Confirmation measure	Retained result
Distinct modeled states	24
Complete plans	8
Priority deferrals	2
Evidence holds	3
Finite-model impossibilities	11
Schedulable cases	10
Direct maintenance downtime	93 half-hour node slots
Richer workflow downtime	90 half-hour node slots
Additional completion	None
Final gain from retention or mathematical adaptation	None beyond conventional optimization

A worked mechanism

In one confirmation case, patch work could not begin before slot eight. The direct plan drained the node at slot five, creating avoidable unavailable waiting. The optimizer delayed the drain to slot seven. All six tasks still finished at slot eleven, but downtime fell from eleven to nine node slots. This explains why the objective changed without additional completed work. [P1]

Correct unit and attribution

Across ten schedulable confirmation cases, three saved half-hour node slots equal 1.5 modeled node-hours. They are not measured service hours, repair savings or customer productivity. The experiment does not isolate Backpass from optimization and other workflow differences.

Negative mathematical finding

The native EMSM component refused the maintenance use under its gold-domain scope. A separately declared temporal-accounting adaptation improved on endpoint-only checks but tied conventional temporal accounting. Final schedules did not improve beyond the optimizer. Retaining that result avoids attributing an advantage to added complexity.

The fleet investigation: show the arithmetic

Public evidence

The retained NYC snapshot is dated October 30, 2026. It is descriptive administrative evidence, supplemented with year-to-date averages and annual context. The web data and report reflect the same underlying City evidence, not independent sources. [P2]

Critical group	Target	Available	Shortfall
Collection trucks	1,253	1,092	161
Dual bin collection trucks	581	489	92
Sewer and water	1,014	963	51
Forestry	155	133	22
Material hauling	224	207	17
Traffic	481	471	10
Buses	105	96	9

Calculation

For each group, shortfall = max(target - available, 0). Gross shortfall sums positive gaps. Across all critical groups, gross shortfall was 378. Surpluses elsewhere totaled 201, leaving net shortfall 378 - 201 = 177. The two collection groups contributed 161 + 92 = 253, or 253 / 378 = 66.9% of gross shortfall.

Why aggregation matters

An available ambulance cannot automatically substitute for an unavailable collection truck. The net total can therefore conceal operationally relevant group gaps. The public numbers support prioritizing investigation, with actual service demand and substitutes still required before scheduling repairs.

A second reconciliation

The citywide target minus available count was 27,082 - 25,887 = 1,195. Agency targets summed to 26,010, yielding an agency net gap of 123. The difference is 1,072 in target levels plus 123, not evidence of 1,072 uncounted vehicles. Agency and critical-group views overlap and must not be added.

Historical context and authority checks

Annual measure	FY2025	FY2026
In-service rate	87%	87%
Days target achieved	64%	52%
Average age, months	90.0	100.0
Repair personnel	1,301	1,266
Nominal repair expenditure, millions	$327.0	$358.7

What the annual context supports

A steady average can coexist with fewer days reaching the target. These aggregate measures suggest investigation of readiness pressure. They cannot identify the separate causal effects of staffing, parts delays, age or procurement. Fiscal 2026 ends before the October snapshot; annual and daily populations must not be subtracted as if contemporaneous. [P2]

Data rejected during review

An open-data endpoint returned 39 records, all dated July 15, 2022. The study excluded it from current-readiness calculations. Seasonal N/A rows were excluded rather than treated as zero availability. No daily time series or consecutive-shortage measure was reconstructed.

Selected local authority evidence

Thirteen synthetic checks used an unchanged existing authority component. They covered issuance, exact resource scope, commencement, replay, unknown reservations, reconciliation, revocation and conservation of cumulative allocation. All passed in the harness. The harness supplied outcomes and did not execute maintenance. All 135 operational acceptance cases remained NOT_RUN. [P3]

Next operational proof

A shadow pilot should collect work orders, repair states, parts timestamps, skills, bay occupancy, route demand and observed service outcomes. Compare against a frozen conventional process. Advance only after the actual owners can inspect the evidence and the complete action route is qualified.

What changed and where the evidence lives

Issue found	Implemented response
Development overlap in initial experiment	Withdrew the exposure claim and froze distinct confirmation states
Gross and net fleet gaps conflated	Reported 378 gross, 201 surplus and 177 net separately
Separate implementations implied independent science	Disclosed same-author model and evaluation exposure
New V2 interpretation could rewrite old evidence	Retained original experimental results unchanged
Local control tests could imply field admission	Kept 13 synthetic checks separate from 135 unrun operational cases

HSDLG contribution

The project illustrates evidence states, explicit constraints, bounded review and retained authority history working toward one operating application. It also shows a useful limit: the conventional optimizer explained the final schedule advantage. A credible Epistemic OS must preserve that conclusion even when a more expansive claim would be easier to promote.

Sources

[P1] Coordinating Physical Work Through Governed Decisions / Physical Work Research Monograph, sections on model, confirmation, mathematics and authority.
[P2] Where Municipal Fleet Availability Falls Short / Municipal Fleet Readiness Study, sections on concentration, target reconciliation, historical context and methods. Retained NYC Fleet Daily Service Report, October 30, 2026, and Fiscal 2026 Mayor’s Management Report, printed p.481.
[P3] Methods, Review and Human Decisions, plus thirteen-case local authority exercise. All supplied in V2G Physical World OS final archive. Tables transcribed from those documents and arithmetic checked in this expansion.

Disclosure

Underlying source snapshots, calculation records and selected test evidence can be reviewed under an appropriate NDA or contract. No live equipment control, independent field validation or measured service improvement is claimed.

Additional technical documentation, standards, test records and implementation details are available for scoped diligence under an appropriate NDA or contract.