Beyond "More Cars": Risk-Informed Sourcing for Longer LRVs and Re-Optimized Train Consists in Legacy Canadian LRTs

Danial Faraji, P.Eng., MIEAust CPEng
Senior Rail Operations and Maintenance Engineer
WSP Canada
Edmonton, Alberta
Danial.faraji@wsp.com

Paper prepared for the session
NG.2 - Lessons Learned from Next-Generation Transportation Projects
2026 Transportation Association of Canada (TAC) Conference & Exhibition
Winnipeg, Manitoba

Abstract

Legacy light rail transit (LRT) capacity programmes must reconcile fleet renewal with platform, power and depot constraints. Adding vehicles to a train and procuring longer individual light rail vehicles (LRVs) are alternative configurations of the same asset system. This paper develops a procurement decision framework that combines discrete platform-fit screening, passenger-exchange analysis, service-based fleet sizing, discounted lifecycle cost and evidence-based risk review. Published agency examples establish the practical context; a fully specified, constructed example demonstrates the calculations without claiming empirical validation. Across a 100–110 m platform screening envelope, four 25 m vehicles require 105.5 m and three 33 m vehicles require 104.0 m under the stated allowances. Before accounting for cab consolidation, both configurations provide a nominal Passenger Capacity of 624 passengers (PAX) per train. A saloon-equivalent sensitivity using an author-defined 1.8 m effective cab length shows how removing two cab modules can increase the longer-vehicle consist to approximately 643 PAX (+3.0%), while four platform-side doors per LRV still give it fewer train-level doorways and a longer base dwell. For cost sensitivity, assigning 20% of the shorter-LRV unit price to its two-cab package and scaling the remaining 80% with vehicle length implies a 25.6% longer-LRV unit-price premium. Under the illustrative 30-year cash flows, this gives a 6.2% lower net present cost for the longer-vehicle configuration. These cab-space and cab-cost effects are analytical sensitivities, not supplier dimensions or quotations. The contribution is a traceable method for locating such thresholds and translating them into bidder evidence, acceptance criteria and operational monitoring.

Keywords: light rail transit; rolling stock procurement; platform fit; dwell time; fleet sizing; lifecycle cost; uncertainty.

1 Introduction

Fleet replacement offers a rare opportunity to reconsider the unit of capacity in a mature light rail transit (LRT) system. Retaining the established light rail vehicle (LRV) length and adding vehicles per train can be effective, but the resulting investment extends beyond the fleet. Calgary's 2015 four-car CTrain introduction increased the stated capacity from 600 to 800 passengers per train. The accompanying programme included platform extensions at 45 stations, signal changes and new substations.[1]

Longer individual vehicles create a different set of interfaces. Toronto's 2014 streetcar introduction plan associated greater capacity and accessibility with higher power requirements, facility access and route conversion works.[2] Boston's 2022 Type 10 procurement presentation distinguished legacy-fleet replacement using single vehicles from a later capacity stage requiring infrastructure for two-vehicle operation.[3] These are documented programme examples, not controlled comparisons of the costs or performance of alternative vehicle lengths.

The procurement question is therefore conditional: under what infrastructure, passenger-flow and cost assumptions can fewer, longer LRVs provide at least the required service at acceptable whole-life cost and risk? This paper addresses that question through a screening framework and a reproducible numerical demonstration. The framework is intended for option development and request-for-proposals preparation. It does not replace detailed design, safety assurance or an agency business case.

2 Background and research contribution

The Transit Capacity and Quality of Service Manual treats rail and station capacity through interacting service, vehicle and infrastructure variables. Passenger Capacity per train alone cannot establish the number of PAX a line can accommodate. The controlling headway, passenger exchange and operational margins also matter.

Published research provides a stronger basis for dwell analysis than a vehicle-length proxy. D'Acierno et al. model dwell as dependent on passenger flows and examine the interaction between crowding and service performance in a metro application.[4] Seriani and Fujiyama's laboratory experiments examine door width, vestibule setback and vertical gap. Their boarding and alighting findings are not identical, illustrating why a nominal increase in door width cannot be treated as a universal dwell improvement.[5] Neither study supplies calibrated Canadian LRT parameters for the example developed here.

The American Public Transportation Association (APTA) Modern Streetcar Vehicle Guideline describes interdependencies among vehicle configuration, platform interface, track interface and power supply.[6] Its stated scope concerns modern low-floor streetcars; high-floor and heritage vehicles fall outside that scope. It is used here for relevant interface principles, not as a universal compliance standard for

[1] City of Calgary (2015), fact sheet. See Reference 1.

[2] Toronto Transit Commission (2014), slides 17, 24–25. See Reference 2.

[3] Peña and Wolfgang, MBTA (2022), slides 2–6. See Reference 3.

[4] D'Acierno et al. (2017), Urban Rail Transit 3: 73–89. See Reference 5.

[5] Seriani and Fujiyama (2019), Collective Dynamics 4, A22. See Reference 6.

[6] APTA (2013), Modern Streetcar Vehicle Guideline. See Reference 7.

Canadian LRT. The lifecycle and interface perspective is also consistent with the INCOSE Systems Engineering Handbook.[7]

The contribution of this paper is the explicit connection of consist-count thresholds, service-equivalent fleet quantities, passenger-exchange assumptions and lifecycle break-even conditions in a procurement workflow. These constituent methods are established engineering tools. The proposed contribution is their integration and traceability, rather than a new capacity law or a claim that previous research has ignored system interfaces. The literature review is targeted to those relationships and is not a systematic review.

3 Problem definition and study design

Three pathways are retained from the procurement problem: A, renew the established vehicle size and operate more vehicles per train; B, procure longer individual vehicles and operate fewer per train; and C, phase the change through a mixed fleet. The numerical comparison evaluates A and B at the same normal-service headway and passenger load policy. C remains a transition pathway requiring a delivery sequence, retained-fleet condition and interim service plan before it can be costed fairly.

The study has three steps. First, public agency documents and relevant technical literature identify the interfaces that require evidence. Second, deterministic screening equations compare configuration feasibility, offered capacity and fleet cost. Third, one-at-a-time and combined sensitivities locate changes in ranking or feasibility. All numerical inputs in the worked example are author-defined analytical assumptions. They are not observations from an operating railway, supplier quotations, fitted estimates, or values taken from non-public employer, client or procurement records. No probability distributions, confidence intervals or claimed probabilities of success are assigned.

The comparison boundary is a complete replacement fleet for one hypothetical line over 30 years. Both options deliver the same normal timetable. Existing fleet purchase costs are not assigned a zero value to favour A, and the comparison does not combine the cost of a full B fleet with only the incremental vehicles needed for A. Passenger time benefits, fare revenue, financing, taxes and wider economic effects are outside the operator cost boundary. Infrastructure changes and major fleet renewals are inside it.

4 Methodology

Figure 1 organizes the decision around mandatory feasibility gates. Unresolved safety, accessibility or interface requirements prevent an option from entering the final trade-off; a favourable price cannot compensate for failure. Service and cost calculations are then carried out together, because headway, fleet quantity and infrastructure requirements are linked.

[7] INCOSE (2023), Systems Engineering Handbook, fifth edition. See Reference 8.

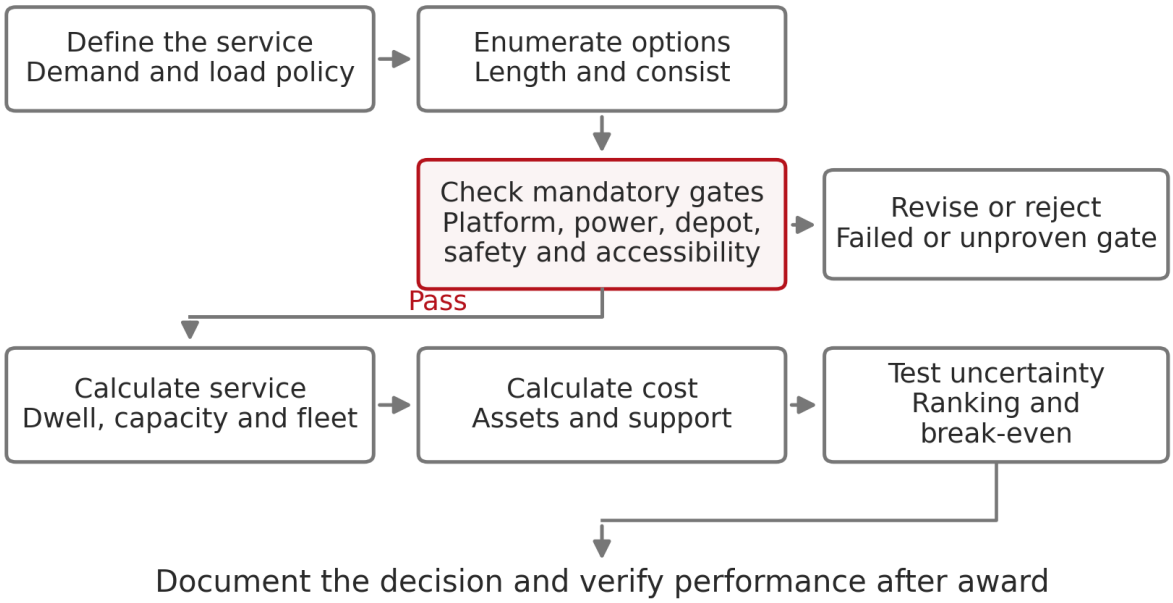


Figure 1. Proposed procurement decision framework. Source: Author. Safety and compatibility gates precede comparative scoring.

4.1 Platform fit and practical train capacity

For identical vehicles and a straight-platform screening envelope, the length required by a consist is

$$L_{\mathrm{req}} = cL_{\mathrm{v}} + (c-1)g + 2m \leq L_{\mathrm{p}} \tag{1}$$

where c is the integer number of vehicles, L_{v} is the vehicle length in metres, g is the additional separation between adjacent vehicle reference envelopes, m is the allowance at each platform end, and L_{p} is usable platform length. The length convention must be explicit: if the supplier's coupled-length definition already includes the inter-vehicle allowance, g is set to zero to prevent double counting. The maximum number of vehicles that can fit is

$$c_{\max} = \max\left\{0, \operatorname{floor}\left(\frac{L_{\mathrm{p}} - 2m + g}{L_{\mathrm{v}} + g}\right)\right\} \tag{2}$$

All operating platforms, turnback locations and other stopping constraints must be screened. Equation (2) is a necessary geometric check, not proof of safe door operation. Curvature, stopping accuracy, coupler movement, accessibility boarding points, sighting, emergency egress and dynamic clearances require separate evidence. Selective door operation is an alternative operational design requiring explicit approval and validation; it is not assumed in the example.

Practical train capacity follows the specified interior and load policy:

$$Q_{\mathrm{train}} = c(n_{\mathrm{s}} + \rho A_{\mathrm{f}}) \tag{3}$$

Here, n_{s} is seats per vehicle, A_{f} is usable standing area in square metres, and ρ is the adopted standing density in passengers per square metre. Door approaches, inaccessible floor area and reserved spaces must be treated consistently. The example uses a hypothetical density of 4 passengers/m²; it does not

prescribe a Canadian load standard. Structural and traction design loads remain separate engineering checks.

4.2 Dwell and service headway

For a doorway at which alighting precedes boarding, a transparent station dwell approximation is

$$t_{\mathrm{d}} = t_{\mathrm{o}} + t_{\mathrm{c}} + t_{\mathrm{buf}} + \max_{\mathrm{d}}\left(\frac{A_{\mathrm{d}}}{q_{\mathrm{a,d}}} + \frac{B_{\mathrm{d}}}{q_{\mathrm{b,d}}} + t_{\mathrm{acc,d}}\right) \quad (4)$$

The opening, closing and dispatch-buffer times are denoted by the first three terms. At doorway d, alighters and boarders are A and B with the doorway subscript, the corresponding flow rates are q_{a} and q_{b} in passengers per second, and the final term covers additional access-related time. The largest doorway exchange controls dwell. Taking only the larger of the separate alighting and boarding times would assume parallel processes; it is not appropriate for the sequential movement assumed here. Opposing flows and congestion may require a calibrated model beyond this approximation.

Where doorway counts are unavailable, the numerical demonstration uses a common-distribution approximation:

$$t_{\mathrm{d}} = t_{\mathrm{fix}} + \frac{1}{\eta D}\left(\frac{A}{q_{\mathrm{a}}} + \frac{B}{q_{\mathrm{b}}}\right) \quad (5)$$

A and B now mean total alighting and boarding passengers for the train at the station. D is the number of available doorways on the platform side, and η is a dimensionless utilization factor in the interval greater than zero and up to one. Thus ηD represents effective doorway capacity. The fixed allowance includes door movement, dispatch buffer and the stipulated access allowance. This simplification assumes comparable flow rates and similar spatial concentration of boarding and alighting. It must be replaced by doorway observations where these conditions do not hold. No benefit is assigned to vehicle length itself.

A screening relationship links dwell to offered line capacity:

$$H = \max(H_{\mathrm{sch}}, H_{\mathrm{other}}, t_{\mathrm{clear}} + t_{\mathrm{d}}), \quad Q_{\mathrm{line}} = \frac{3600}{H} Q_{\mathrm{train}} \quad (6)$$

is the scheduled headway, represents other controlling constraints, and the clearance term excludes the dwell already included. All times are in seconds, and is offered Passenger Capacity in PAX per hour per direction. The headway expression is a planning surrogate, not a signalling design formula. A project must establish constraints for its junctions, terminals, traffic signals and train control. Offered Passenger Capacity is distinct from delivered or realized throughput; no additional loading or reliability multiplier is applied to duplicate the load policy in Equation (3).

4.3 Fleet quantities and lifecycle cost

For a common normal-service cycle-time budget, the active train and purchased vehicle quantities are

$$N_{\mathrm{tr}} = \mathrm{ceil}\left(\frac{T_{\mathrm{cyc}}}{H_{\mathrm{sch}}}\right), \quad N_{\mathrm{fleet}} = \mathrm{ceil}((1 + s)cN_{\mathrm{tr}}) \quad (7)$$

T_{cyc} includes the round trip, scheduled station dwells and terminal recovery. The spare ratio s is applied to the active vehicle requirement, with rounding upwards. This is a planning allowance, not a fleet

availability model. Major changes in dwell, headway or turnaround require the timetable and fleet quantities to be recalculated together.

Net present cost (NPC) is calculated in a consistent real price base:

$$\mathrm{NPC_j} = C_{\mathrm{fleet,j,0}} + C_{\mathrm{inf,j,0}} + \sum_{t=1}^{T} \frac{C_{\mathrm{j,t}}}{(1+r)^t} - \frac{S_{\mathrm{j,T}}}{(1+r)^T} \quad (8)$$

The initial terms represent fleet and infrastructure capital at year zero for option j. The annual cash flow includes operations and maintenance (O&M), energy, major renewals, later infrastructure works and any separately estimated monetary risk consequences in year t. The terminal residual value is S, the real discount rate is r, and the analysis period is T. Monetary risks must be entered once: an expected loss already included in the cash flow is not added again as an identical contingency or procurement-score penalty. Unpriced risks remain visible in the risk register.

The worked example adopts a 30-year horizon and a 4% real discount rate as analytical assumptions. Neither value is presented as a universal Canadian requirement. All figures are constant 2026 Canadian dollars, with end-of-year recurring payments and a year-15 renewal. Price escalation is not added to real cash flows. Agency-specific discounting, useful life and residual-value policies would replace these assumptions in an investment appraisal.

4.4 Risk and procurement scoring

Only options that satisfy the mandatory gates proceed to comparative scoring. If a weighted score is used, each criterion must have an agreed value function and published anchors:

$$S_j = \sum_{k=1}^{K} w_k v_{j,k}, \quad \sum_{k=1}^{K} w_k = 1 \quad (9)$$

The weights are non-negative and the criterion values are normalized to a common 0–100 scale. An assessor's ordinal likelihood and severity ratings are not multiplied and subtracted from this score as though they were comparable monetary or interval-scale quantities. Evidence maturity is instead recorded explicitly, unresolved mandatory evidence is a gate, and sensitivities test disputed assumptions. No numerical bid score is claimed in this study because no bidder evidence has been evaluated.

5 Illustrative application and results

The numerical application considers a hypothetical line moving from three shorter vehicles per train to either A, four shorter vehicles, or B, three longer vehicles. Table 1 specifies the operating assumptions. The geometry, Passenger Capacity, passenger-flow and timetable values are deliberately explicit so that the consequences of changing them can be inspected. They are author-defined for this methodological demonstration and do not describe a marketed vehicle, a named Canadian line, an unpublished procurement or an internal project.

Table 1. Author-defined operating assumptions for the numerical demonstration.

Input	A Shorter vehicles	B Longer vehicles
Vehicles per train; vehicle length	4; 25 m	3; 33 m
Seats; usable standing area per vehicle	48; 27 m²	64; 36 m²

Input	A Shorter vehicles	B Longer vehicles
Doorways per vehicle on platform side	4	4
Door utilization factor	0.70	0.80
Usable platform screening envelope; allowance at each end	100–110 m; 2 m	Same
Additional inter-vehicle gap	0.5 m	Same
Standing density	4 passengers/m²	Same
Boarding; alighting passengers per train	240; 160	Same
Boarding; alighting rate per doorway	1.0; 1.3 passengers/s	Same
Fixed dwell allowance	20 s	Same
Scheduled; other minimum headway	180 s; 150 s	Same
Station clearance excluding dwell	110 s	Same
Cycle-time budget; spare ratio	5,400 s; 15%	Same
Cabs per bidirectional LRV; effective cab length sensitivity	2; 1.8 m/cab	Same

The 20 s fixed allowance comprises 10 s for door operation, 5 s for the stipulated access allowance and 5 s for the dispatch buffer. It is not a measured mean or a percentile. The 90 min cycle is a common timetable budget, with sufficient scheduled dwell and recovery for the two base configurations; route running times have not been simulated. Doorway flow rates incorporate their assumed effective width, so doorway count is not multiplied by width a second time.

5.1 Geometry and equal normal-service capacity

Table 2. Calculated base results from Table 1.

Output	A	B
Length including allowances	105.5 m	104.0 m
Remaining length at 110 m worked point	4.5 m	6.0 m
Nominal Passenger Capacity per vehicle	156 PAX	208 PAX
Nominal Passenger Capacity per train	624 PAX	624 PAX
Total cabs per consist	8	6

Output	A	B
Saloon-equivalent length at 1.8 m/cab	85.6 m	88.2 m
Cab-adjusted Passenger Capacity sensitivity	624 PAX	643 PAX
Physical; effective available doorways	16; 11.2	12; 9.6
Calculated dwell	52.4 s	57.8 s
Screened service headway	180 s	180 s
Cab-adjusted offered Passenger Capacity	12,480 PAX/h	12,859 PAX/h
Active trains; purchased vehicles	30; 138	30; 104

At the same load policy and 180 s headway, the three-short-vehicle baseline supplies 468 PAX per train and 9,360 PAX/h/direction. Option A increases nominal Passenger Capacity to 624 PAX/train. Option B also gives 624 PAX/train from the explicit seat-and-standing calculation before any cab-space adjustment. The comparison is then extended with a separate cab-count sensitivity because four bidirectional LRVs contain eight cabs whereas three bidirectional LRVs contain six.

A four-to-three vehicle change removes two cab modules from the train consist. To expose the possible saloon-space effect without claiming a supplier layout, an effective cab length of 1.8 m per cab is used as the base sensitivity. The saloon-equivalent train length is therefore $4(25 - 2\times1.8) = 85.6$ m for A and $3(33 - 2\times1.8) = 88.2$ m for B. Calibrating the same interior-capacity intensity to A's 624 PAX gives approximately 643 PAX for B, an increase of 19 PAX or 3.0%. A 1.5–2.0 m/cab sensitivity range gives approximately 638–646 PAX for B (+2.3% to +3.6%). The threshold is transparent: because the total nominal vehicle length is 100 m for A and 99 m for B, B has more saloon-equivalent length whenever the effective cab length exceeds 0.5 m. Actual procurement evaluation must use bidder floor plans, cab boundaries, equipment cupboards, articulation zones, accessible spaces and seating layouts rather than this proxy.

The platform check is expressed as a 100–110 m screening envelope rather than as a single nominal platform. Under the stated 2 m end allowances and 0.5 m additional inter-vehicle gaps, A requires 105.5 m and B requires 104.0 m. Neither option fits below 104.0 m; only B fits from 104.0 m to less than 105.5 m; and both fit from 105.5 m to 110 m. Exact equality at a threshold leaves no additional tolerance beyond the stipulated allowances and requires route-specific verification.

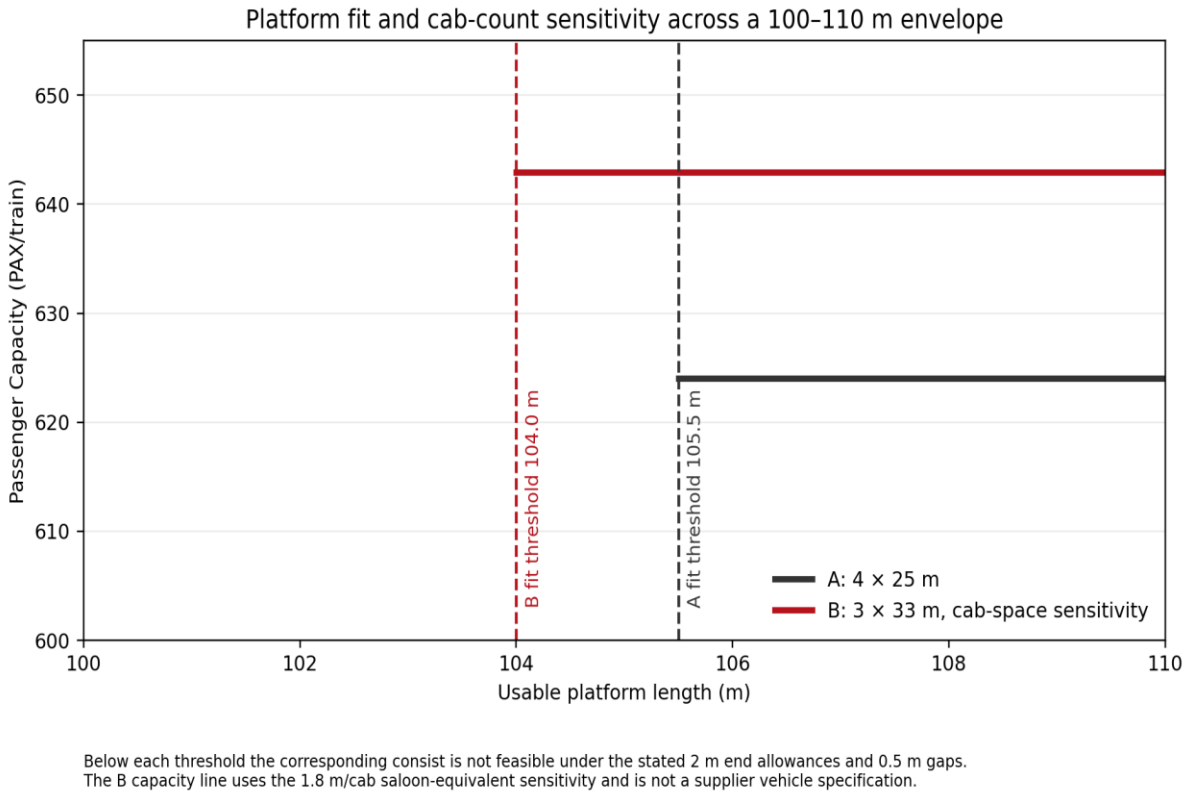


Figure 2. Platform-fit and cab-count sensitivity across a 100–110 m usable-platform screening envelope. Option A is 624 PAX/train; Option B is approximately 643 PAX/train only under the author-defined 1.8 m/cab saloon-equivalent sensitivity. Source: Author calculations.

5.2 Door performance and adverse operation

With four doorways per side on both vehicle types, B has 12 physical platform-side doorways per train, compared with 16 for A. At the base utilization factors, A therefore has 11.2 effective doorways and B has 9.6. The calculated base dwell is 52.4 s for A and 57.8 s for B. The condition for B to have no worse aggregate passenger-exchange time, at equal flow rates and PAX totals, is that its effective doorway count is at least A's. With 12 versus 16 physical doorways and A's utilization of 0.70, B would require a utilization factor of at least 0.933. The cab-space capacity sensitivity therefore does not imply a dwell benefit; doorway count, utilization and PAX distribution remain separate design variables.

Table 3. Conditional operating sensitivities using the screening headway model.

Scenario	**Dwell A (s)**	**Dwell B (s)**	PAX/h A	PAX/h B
Base	52.4	57.8	12,480	12,859
Both utilization factors 0.70	52.4	63.2	12,480	12,859
B utilization reduced to 0.65	52.4	66.5	12,480	12,859

Scenario	Dwell A (s)	Dwell B (s)	PAX/h A	PAX/h B
One doorway unavailable per train	54.6	61.3	12,480	12,859
Flow rates reduced by 15%	58.1	64.5	12,480	12,859
Exchange × 1.75 and flow rates −15%	86.7	97.9	11,418	11,135

The 15% flow reduction is an assumed adverse-flow stress, potentially relevant to bulky clothing or constrained circulation; it is not a measured winter penalty. The combined case has 420 boarding PAX and 280 alighting PAX per train. It increases the screened headways to 196.7 s for A and 207.9 s for B. Using the cab-adjusted Passenger Capacity sensitivity, offered line capacity becomes approximately 11,418 PAX/h/direction for A and 11,135 PAX/h/direction for B. Passenger totals are held fixed within each sensitivity to isolate the doorway mechanism. Further PAX accumulation, dwell feedback, running-time changes and recovery would require a network simulation. The adverse-case outputs are therefore conditional screening results, not equilibrium forecasts.

A door failure is modelled by removing one available doorway without changing interior capacity. Continued operation with that door isolated would require confirmation against the applicable operating and emergency-egress rules. Separate checks are needed for an inaccessible boarding position or a failure requiring vehicle withdrawal.

5.3 Lifecycle cost and break-even conditions

Table 4 defines the cash flows for the complete replacement fleets. All monetary inputs are author-defined scenario values in constant 2026 CAD and are not agency estimates, supplier quotations or values extracted from a live procurement. The infrastructure allowances cover configuration-specific station, power, depot and integration work in aggregate. Their difference is stipulated to demonstrate the trade-off and is not inferred from the vehicle-length difference. O&M excludes energy, and major renewals are excluded from both annual categories to avoid double counting. Both options retain the same train staffing assumption; no driver saving is assigned simply because there are fewer vehicles. To test the cab-cost hypothesis transparently, the shorter-LRV unit price is decomposed into a 20% fixed two-cab package and an 80% remainder that scales with vehicle length. The implied longer-LRV price ratio is therefore 0.20 + 0.80(33/25) = 1.256, corresponding to CAD 7.536 million per LRV, or a 25.6% unit-price premium. At consist level, A has a purchase-cost index of 4.000 short-LRV equivalents and B has 3×1.256 = 3.768, 5.8% lower. The two fewer cab modules account for 0.20 of one short-LRV price per consist within this decomposition. This is an analytical sensitivity only: once an actual bidder price is available, the 20% cab share must not be deducted again.

Table 4. Illustrative cash-flow inputs in million constant 2026 CAD.

Input and timing	A	B
Vehicle unit price at year 0 (derived sensitivity)	6.000	7.536

Input and timing	A	B
Purchased vehicles from Equation (7)	138	104
Calculated fleet capital at year 0	828.0	783.7
Infrastructure allowance at year 0	120.0	85.0
Annual O&M, years 1–30	24.0	23.0
Annual energy, years 1–30	6.0	5.8
Major renewal at year 15	138.0	145.0
Residual value at year 30	0.0	0.0
Cab-related share of shorter-LRV unit price	20%	Same analytical basis
Derived B unit-price premium relative to A	—	25.6%
Real discount rate; analysis period	4%; 30 years	Same

The common recurring-cost annuity factor is 17.2920. With the derived CAD 7.536 million B unit price, calculated NPC is CAD 1,543.4 million for A and CAD 1,447.3 million for B (Figure 3). B's CAD 96.1 million advantage is 6.2% of A's cost. CAD 35.0 million of this difference is the stipulated initial infrastructure allowance; the remaining difference is dominated by the lower complete-fleet purchase cost created by the three-vehicle consist and the assumed unit-price relationship. The result is therefore a sensitivity to the cost decomposition, not evidence that every longer LRV procurement will achieve this saving.

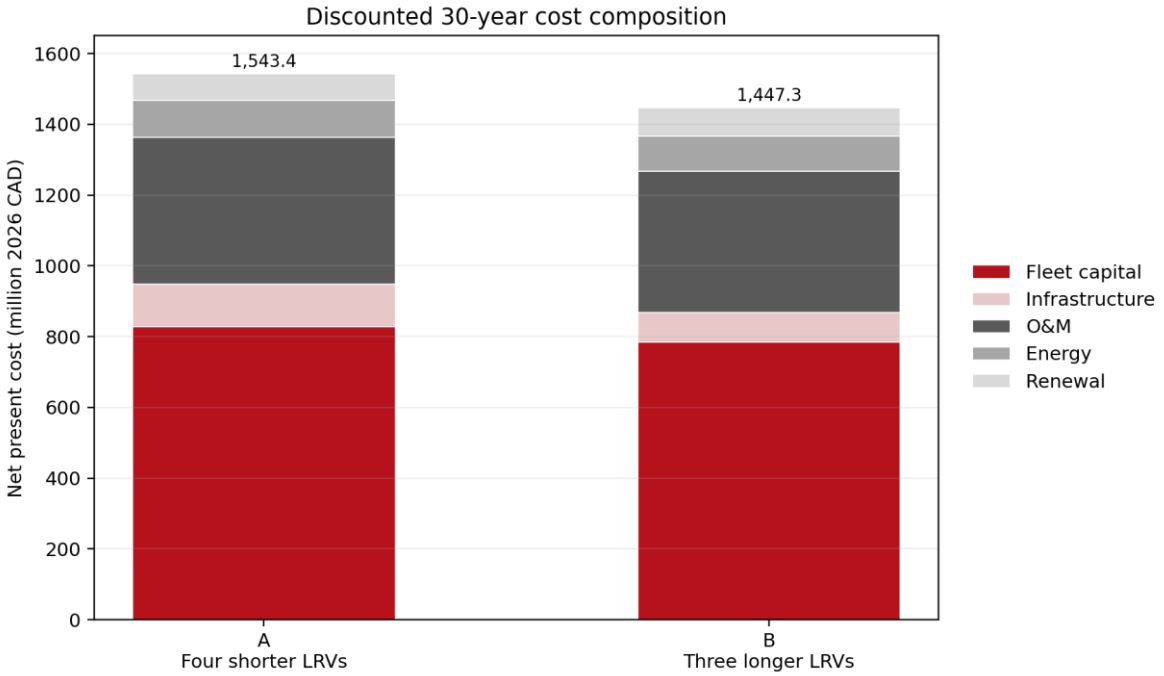


Figure 3. Discounted cost composition over 30 years at a 4% real rate, using the 25.6% derived longer-LRV unit-price premium. All inputs are illustrative and totals are in million constant 2026 CAD. Source: Author calculations from Table 4.

Holding the other Table 4 inputs constant, B can absorb approximately CAD 96.1 million of additional upfront cost before matching A. Allocating that amount across 104 vehicles gives a break-even B unit price of CAD 8.460 million, approximately 12.3% above the derived CAD 7.536 million base price and about 41.0% above the CAD 6.0 million short-LRV price (Figure 4). The stated 20–30% longer-LRV unit-price premium range therefore remains below the break-even value in this constructed example. This is a scenario boundary, not a market estimate or an acceptable bid price for a real procurement.

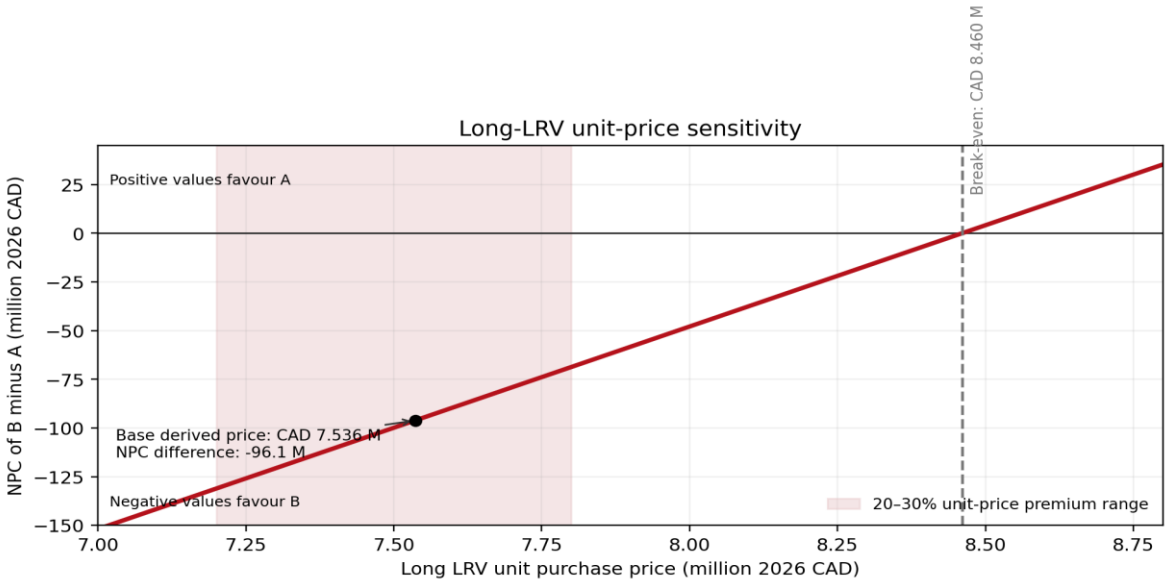


Figure 4. Sensitivity to the longer-LRV unit purchase price. The 20–30% premium range is shown against the CAD 8.460 million break-even value, with all other inputs fixed. Source: Author calculations.

Table 5. Cost sensitivities in million constant 2026 CAD.

Scenario	NPC A	NPC B	B minus A
Base: 25.6% B unit premium	1,543.4	1,447.3	-96.1
B unit-price premium 20%	1,543.4	1,412.3	-131.1
B unit-price premium 30%	1,543.4	1,474.7	-68.7
Equal annual recurring costs	1,543.4	1,468.0	-75.4
No infrastructure saving for B	1,543.4	1,482.3	-61.1
Infrastructure A -25%, B +25%	1,513.4	1,468.5	-44.9

Across the stated 20–30% longer-LRV unit-price premium range, B remains lower in NPC in the constructed example: approximately CAD 1,412.3 million at a 20% premium and CAD 1,474.7 million at a 30% premium, compared with CAD 1,543.4 million for A. Equal annual recurring costs reduce B's advantage to CAD 75.4 million. Removing the assumed infrastructure saving still leaves a CAD 61.1 million advantage, while the adverse infrastructure case (A −25%, B +25%) leaves a CAD 44.9 million advantage. These are deterministic scenario boundaries, not confidence bounds. The cost ranking remains highly dependent on actual supplier pricing, facility scope, renewals and support obligations.

6 Structured risk register

A risk register connects the calculation assumptions to evidence and responsibility. Table 6 proposes the minimum topics for this decision. Each live register entry should also record an owner, affected requirement, current evidence, proposed control, residual risk, due date and acceptance authority. The entries below identify what needs to be established; they do not assign project probabilities or certify any configuration.

Table 6. Proposed interface risks and evidence required before commitment.

Risk mechanism	Evidence or control	Acceptance basis
Platform or swept-envelope mismatch	Platform-by-platform length and door-location matrix; curve and stopping-tolerance analysis.	Verified against each operating location.
Dwell and access performance below target	Doorway observations or validated passenger simulation; accessibility boarding trials.	Target dwell and passenger distribution at controlling stations.
Traction power or braking incompatibility	Loaded-train power and braking studies, simultaneous departures and degraded cases.	Network and safety requirements met under stated load cases.
Depot and rescue mismatch	Bay, lift, wheel-lathe, wash and roof-access drawings; rescue and failed-train plan.	Facilities and operating procedures demonstrated for the proposed vehicle.
Winter service degradation	Door, heating, ventilation and air-conditioning, snow and ice tests; maintenance access assessment.	Defined winter duty and recovery criteria, with evidence.
Mixed-fleet and support exposure	Configuration, spares, training, software support and introduction plans.	Transition service maintained and support obligations traceable.

Several risks are correlated. Snow conditions can affect both door availability and passenger flow; a vehicle interface change can affect the depot, rescue plan and training simultaneously. Multiplying unrelated-looking risk scores would not resolve those dependencies. Combined scenarios should be selected around credible common causes, and material cost consequences should be added once to the lifecycle model when evidence becomes available. The worked cost example excludes unestimated risk allowances rather than implying that risk has no cost.

7 Procurement playbook

7.1 Specify outcomes and disclose the interface baseline

The request for proposals should state the required offered service, passenger load policy, critical-station dwell criteria, platform envelopes, accessibility requirements, winter duty, facility limits and degraded operating modes. Agencies should provide the same infrastructure drawings, passenger data and uncertainty statements to each proponent. Vehicle length can then be offered as a design variable within an explicit feasible range, with evidence required for every proposed departure.

7.2 Require configuration-specific evidence

Each proponent should submit an infrastructure change register distinguishing mandatory works, risk-reduction measures and optional enhancements. The vehicle layout, door map, coupled-length definition, traction and braking assumptions, maintenance provisions and rescue interfaces must correspond to the configuration being priced. A generic product family brochure is insufficient evidence for a route-specific claim. The Federal Transit Administration's Rail Vehicle Qualification Test Compendium provides a reference for organizing qualification evidence; the applicable project requirements and acceptance programme still govern.[8]

7.3 Separate mandatory gates from comparative value

An illustrative weighting for compliant options is 30% lifecycle cost, 20% practical capacity, 15% dwell performance, 15% transition delivery, 10% long-term support and 10% maintainability and operational resilience beyond the mandatory minimum. These are proposed weights, not results of stakeholder elicitation. Agencies should set criterion anchors and test the effect of weight changes before viewing bids. Correlated advantages, such as a maintenance saving included in both NPC and a qualitative score, require an explicit rationale to avoid rewarding the same outcome twice.

7.4 Verify the selected configuration in service

The contract should trace every material assumption to a verification activity and acceptance authority. Monitor dwell distributions by station and passenger load, unavailable doors, load distribution, energy per train-kilometre and passenger-kilometre, winter service-affecting failures, maintenance labour and delivered train availability. The cost and performance baseline should be revisited when infrastructure scope, vehicle layout or operating rules change. Post-award monitoring is part of closing the decision, not proof that the original screening model was validated.

8 Comparative lessons from public examples

Calgary demonstrates that adding a car can be a workable capacity programme when supporting station, signal and power changes are delivered together.[1] Toronto illustrates the need to integrate accessible vehicles with power, storage and route conversion.[2] Boston's published Type 10 staging separates initial replacement from a later two-vehicle capacity increase, with supporting infrastructure explicitly identified.[3]

The transferable lesson is the need to define the service configuration and its dependencies. The examples do not establish a transferable percentage cost saving, common passenger-density standard or causal relationship between vehicle length and dwell. Toronto streetcar and Boston legacy-rail experience can inform the questions asked in Canadian LRT procurement, while local track geometry, platform height, service control and operating practices determine applicability.

9 Discussion and limitations

The example identifies three different decision thresholds. A length threshold can remove an entire vehicle from a feasible consist; a doorway-utilization threshold can reverse the dwell comparison; and a

[8] Federal Transit Administration (2015), Report No. 0083. See Reference 9.

purchase-price or infrastructure threshold can reverse the cost ranking. These thresholds lead to specific procurement questions: what length definition has the bidder used, how will passengers distribute across doors, and which asset changes are included in the price?

Fewer vehicles are not equivalent to less equipment or lower maintenance burden. The number of articulated modules, bogies, traction packages, doors and maintainable components may change. The cab-count reduction can release saloon-equivalent space and avoid two cab modules at consist level, but a longer unit also concentrates more PAX capacity in each unavailable vehicle. Spare strategy, repair duration, train reformation, coupling compatibility and partial service operation therefore need configuration-specific assessment. Equations (7) and (8) do not establish a reliability benefit from the lower vehicle count.

The main evidence limitation is the absence of calibrated local measurements and comparable supplier quotations. The numerical results establish the mathematical consequences of declared assumptions, not the prevalence of those conditions in Canadian networks. The door-utilization factors, flow rates, cost allowances and common cycle time have not been calibrated. The geometric check does not demonstrate clearance, traction-power or maintenance feasibility, and the headway surrogate does not reproduce junction conflicts, traffic signal delay or terminal recovery. The current work therefore supports screening, requirement definition and evidence planning rather than empirical validation.

A field application should obtain doorway-level passenger counts, actual layouts, station geometry, timetable and power data, maintenance task times and comparable price submissions. Passenger-flow and service models should then be calibrated on one dataset and assessed against separate operating observations, including winter and disrupted service. A detailed fleet availability assessment and time-phased model are needed for the mixed-fleet pathway. These extensions would permit testing whether the identified thresholds predict real operational or procurement outcomes.

The practical value for TAC is a decision process that can make hidden infrastructure scope and untested operating assumptions visible before contract award. The integration is useful even when no option dominates: an unresolved ranking can direct the next survey, supplier clarification or design study. The framework should preserve that uncertainty rather than convert incomplete evidence into an apparently precise score.

10 Conclusion

LRV length and consist policy should be procured as a coupled fleet, infrastructure and operating configuration. The proposed framework first screens mandatory interfaces, then connects practical capacity and dwell assumptions to headway, fleet quantity and lifecycle cost. Its contribution is a traceable comparison and a set of testable decision thresholds.

In the constructed example, four 25 m vehicles and three 33 m vehicles each provide a nominal 624 PAX/train before cab-space adjustment. Because the three-vehicle consist contains two fewer cab modules, the 1.8 m/cab saloon-equivalent sensitivity increases B to approximately 643 PAX/train (+3.0%), while its lower train-level doorway count produces a longer base dwell. A 20% cab-related cost share combined with length-scaling of the remaining unit cost implies a 25.6% longer-LRV unit-price premium and a 6.2% lower NPC for B under the stated cash flows. These results do not establish a universal preference for longer vehicles; they show why cab count, usable saloon space, door distribution, supplier price and infrastructure scope should be evaluated together and verified from bidder-specific evidence.

Data and calculation availability

The numerical study is based entirely on public sources and author-defined analytical inputs. Tables 1 and 4 state the principal calculation inputs, and the accompanying reproducibility workbook contains only those stated inputs, transparent formulas and sensitivity calculations. No non-public employer, client, supplier or procurement document is used as a data source. No stochastic sampling is used.

References

1. City of Calgary. “Calgary Transit launches four-car train service early.” News release and fact sheet, 13 November 2015. Available online. Accessed 17 September 2026.

2. Toronto Transit Commission. “Low Floor Streetcar Key Features and Service Launch Preparedness.” TTC Board presentation, 19 August 2014, slides 17, 24–25. Available online. Accessed 17 September 2026.

3. Peña, A., and Wolfgang, B. “RFP No 367F-19 Type 10 Supercar Update and Procurement Award.” Massachusetts Bay Transportation Authority, 31 August 2022, slides 2–6. Available online. Accessed 17 September 2026.

4. Transportation Research Board. Transit Capacity and Quality of Service Manual, Third Edition. TCRP Report 165. Washington, DC: The National Academies Press, 2013. Rail and station capacity chapters. doi: 10.17226/24766. Accessed 17 September 2026.

5. D’Acierno, L., Botte, M., Placido, A., Caropreso, C., and Montella, B. “Methodology for Determining Dwell Times Consistent with Passenger Flows in the Case of Metro Services.” Urban Rail Transit 3 (2017): 73–89. doi: 10.1007/s40864-017-0062-4.

6. Seriani, S., and Fujiyama, T. “Exploring the Effect of Train Design Features on the Boarding and Alighting Time by Laboratory Experiments.” Collective Dynamics 4 (2019): A22, 1–21. doi: 10.17815/CD.2019.22.

7. American Public Transportation Association. Modern Streetcar Vehicle Guideline. APTA RT-ST-GL-001-13, approved 26 March 2013. Introduction and Chapters 1–4. Available online. Accessed 17 September 2026.

8. INCOSE. Systems Engineering Handbook: A Guide for System Life Cycle Processes and Activities. Fifth edition. Hoboken, NJ: Wiley, 2023. ISBN 978-1-119-81429-0. Available online.

9. Federal Transit Administration. Rail Vehicle Qualification Test Compendium. FTA Report No. 0083. Washington, DC: U.S. Department of Transportation, February 2015. Available online. Accessed 17 September 2026.