Kingston's ALTO Ridership Analysis

Today's service, an upgraded conventional railway, and two ALTO station locations

Policy Brief · July 2026

ALTO's chief executive has said Kingston will probably receive a station, and that most trains would pass through without stopping. Neither the service pattern nor the station location has been published. It compares today's railway, an upgraded conventional railway at the existing station, and ALTO with a station either in the city or twenty-seven minutes north, at conventional and premium fares. Every parameter is set out in full so that any can be challenged.

SUMMARY OF FINDINGS

- Only one option raises Kingston's rail ridership: an upgraded conventional railway at the existing station, about +12 per cent. The best ALTO case — an in-town station at conventional fares — merely matches today. Every other ALTO case falls below it, by 8 to 17 per cent.
- ALTO's journey-time advantage at Kingston is worth about +10 per cent on its own. Reducing calls from eighteen to eight gives back the same amount; a 25 per cent fare premium gives it back again; moving the station out of the city costs six to eight points more.
- Frequency cannot repair this on its own. Even at eighteen calls a day — parity with today — an out-of-town station at a premium fare still returns about −9 per cent.
- Kingston's intra-corridor trips — Belleville, Brockville, Cobourg, Napanee, Oshawa, about 8 per cent of volume — have no ALTO equivalent at any frequency.

TABLE 1 · HEADLINE COMPARISON

Option	Toronto	Calls a day	Premium	Annual trips	Change
Current service	135 min	18	none	450,000	—
HPPR, existing station	95 min	18	none	502,000	+12%
ALTO, in-town station	80 min	8	none	449,000	0%
ALTO, in-town station	80 min	8	+25%	403,000	−10%
ALTO, 27 min north	80 min	8	none	416,000	−8%
ALTO, 27 min north	80 min	8	+25%	376,000	−17%

All four ALTO rows assume eight calls a day and conventional service withdrawn; they differ only in station location and fare. No fare structure for intermediate stations has been published, so both cases are shown rather than assumed. HPPR is the 240 km/h new-build railway proposed under the High Performance Rail framework. Ranges and sensitivities are on pages four and five.

1 · Why Kingston's VIA ridership is already strong

Kingston's standing among the busiest stations on the network is routinely cited as evidence for a high-speed stop. What produces that volume determines whether a different kind of station would reproduce it.

Frequency, from two service patterns

Kingston sits at the midpoint of the Toronto–Montréal mainline and on the path of the Toronto–Ottawa trains, which use the same subdivision as far as Brockville. It therefore collects two service patterns rather than one, giving it a level of service exceeded only by the three largest cities on the corridor. Intercity frequency elasticities typically fall between +0.4 and +0.7, so service density accounts for a substantial share of the volume independently of speed.

The station also serves a catchment wider than the city: Napanee, Gananoque, Amherstview and the western Thousand Islands have no intercity rail of their own and drive to Kingston to board. Volume attributed to a city of 132,000 is generated by a region several times that size.

A population unusually disposed to rail

Roughly forty thousand post-secondary students live in a city of about 132,000 — one of the highest ratios in the country — drawn disproportionately from the Toronto and Ottawa regions and travelling without cars. Kingston also has a large retired population, for whom avoiding the highway is the point of the trip, and an exceptionally high share of institutional and public-sector employment whose travel is expensed and defaults to rail.

This travel is not evenly spread. It concentrates at term boundaries, on Thursday and Sunday afternoons and around reading weeks and holidays, and it generates a second flow in the opposite direction as families travel to Kingston rather than from it. These segments are the most sensitive to frequency and to how far the station sits from where they are going, and the least able to substitute a car. They are also low-yield: peaked at term boundaries and weekends, price-sensitive, and concentrated outside the weekday business hours on which a high-speed operation's revenue depends. Kingston's ridership is strong for reasons that make it hard to monetise at a premium fare.

What the volume is, and is not, evidence of

Kingston's standing as one of the busiest stations rests on statements by the operator and the Minister, not on published station-level data — the first item in the verification list at the end of this brief. It is also a measure of existing demand: it establishes that Kingston already travels by train, not that a different station would generate additional trips. Only generation adds ridership to the corridor.

Distance, and the absence of an air alternative

Toronto is about 260 km away and Ottawa about 170 km — far enough that driving is a chore, close enough that rail's fixed overheads remain tolerable. With no scheduled air service to Toronto, Kingston has no competing mode in its principal market, and this matters for any ridership case built on the city: the premium travellers who elsewhere transfer from air to high-speed rail are already on the train. There is no abstraction pool. Incremental ridership can come only from cars or from induced travel.

Both attributes that produced Kingston's volume — frequency, and a station inside the urban fabric roughly ten minutes from the core and the university — are removed by a high-speed alignment north of the city. That is what the model below is built to test.

2 · Method

The model compares the generalised cost of a Kingston rail journey today with the same journey under each option, and converts the difference into a demand change through a constant-elasticity response. Generalised cost is expressed in minutes: in-vehicle time, access and egress weighted for their greater disutility, a schedule-delay penalty from the interval between trains, any interchange penalty, and the fare converted at each segment's value of time.

Demand is disaggregated into four market pairs and four traveller segments, each split by car availability — thirty-two cells, computed separately and summed, because the question turns on segments responding differently.

TABLE 2 · PARAMETERS

Parameter	Value used
Base volume	450,000 annual station movements (sensitivity band 400,000–550,000)
Market split	Toronto 58%, Ottawa 22%, Montréal 12%, intra-corridor 8%
Segment split	Students 30%, seniors and leisure 25%, public sector 20%, other 25%
Car-free share by segment	Students 85%, seniors and leisure 50%, public sector 15%, other 20%

Parameter	Value used
Value of time	$14, $20, $48, $24 per hour respectively
In-vehicle time	Today 135 / 120 / 160 min; ALTO 80 / 45 / 105 min (Toronto / Ottawa / Montréal)
Kingston-end access	Existing station 10 min by car, 20 by transit; ALTO 27 by car, 35 by shuttle
Access time weight	1.5 where a car is available, 2.0 where it is not
Schedule delay	half the average interval, weighted 0.5, over a fifteen-hour service day
Fare premium, ALTO	25% central; 0% and 40% tested
Generalised cost elasticity	−1.0 central (band −0.8 to −1.2)

Scenarios and policy states

Each service level is tested under two policy states, because the outcome depends less on ALTO than on what happens to the service Kingston already has.

- Substitution. ALTO becomes Kingston's rail service in the three markets it reaches; conventional service is withdrawn or reduced below a useful level. Intra-corridor trips have no rail option.
- Coexistence. Conventional service is retained at present frequency and ALTO added. Travellers take whichever option has the lower generalised cost, and only the improvement over the better of the two generates new demand.

Limits of the estimate

Four inputs are assumptions rather than data, and each appears in the verification list at the end: base volume, market split, ALTO journey times (the alignment for this segment is unpublished) and station location. Far-end access time is held identical throughout, which is conservative. The elasticity is applied to a large change in generalised cost, at the outer edge of where constant-elasticity methods behave well: the direction is robust, the magnitude indicative.

3 · Results

TABLE 3 · WHERE ALTO'S JOURNEY-TIME GAIN GOES

Step	Annual trips	Change	Effect
Current service, as operated	450,000	—	—
Journey time cut to 80 min, all else held	496,000	+10%	+10 pts
Calls cut from 18 to 8 a day	449,000	0%	−10 pts
Fare premium of 25 per cent	403,000	−10%	−10 pts
Station moved 27 minutes north	376,000	−17%	−6 pts

The second row isolates the whole value of high-speed line-haul time at Kingston. It already excludes intra-corridor trips, which a high-speed alignment cannot serve at any frequency.

TABLE 4 · ALTO AT A STATION 27 MINUTES NORTH, BY SERVICE LEVEL

Service level	Substitution	Change	Range	If VIA retained
6 calls a day	359,000	−20%	−17% to −24%	0%
8 calls a day	376,000	−17%	−13% to −21%	0%
10 calls a day	387,000	−14%	−10% to −18%	+0.2%
18 calls a day	408,000	−9%	−4% to −14%	+1.6%
6 per direction (12 total)	394,000	−12%	−9% to −17%	+0.5%
8 per direction (16 total)	404,000	−10%	−6% to −15%	+1.3%

Service level	Substitution	Change	Range	If VIA retained
10 per direction (20 total)	411,000	−9%	−3% to −14%	+1.9%

The range spans elasticity, value of time, market split and segment mix at a 25 per cent fare premium. The lower block reads six, eight and ten as calls per direction — the more generous interpretation. It improves the outcome without changing its sign.

Sensitivity

Two parameters move the result materially. The fare premium is the stronger.

TABLE 5 · CHANGE IN ANNUAL TRIPS, SUBSTITUTION STATE

Scenario	Fare premium 0%	25%	40%	Access 25 min	45 min
6 calls a day	−12%	−20%	−24%	−17%	−23%
8 calls a day	−8%	−17%	−21%	−13%	−20%
10 calls a day	−5%	−14%	−19%	−10%	−17%

Access columns hold the fare premium at 25 per cent and vary the drive time from station to core; the central case is 27 minutes.

Break-even

The more useful question is what would make an ALTO station better for Kingston than the service it has. Holding the station at a twenty-seven-minute drive, the answer is almost entirely about price.

TABLE 6 · DAILY CALLS REQUIRED TO MATCH TODAY'S GENERALISED COST, TORONTO MARKET

Segment	At conventional fares	At a 25% premium
Public sector and institutional	9	13
Other business and leisure	9	27
Seniors, leisure and visiting friends	9	48
Students and young adults	9	not achievable at any frequency

One exception runs the other way: expensed public-sector travel between Kingston and Ottawa, where today's service is slow and indirect, is better off under ALTO in every scenario tested.

Under favourable assumptions on both — conventional fares and a station twenty-five minutes from the core — retaining conventional service alongside ALTO yields about +8 per cent at ten calls a day. That is the best figure the model will produce for an out-of-town Kingston station, and it requires abandoning the premium pricing the revenue case depends on elsewhere.

4 · Speed at constant frequency

The scenarios above vary service while holding journey time at what a high-speed alignment would deliver. This section reverses the test. Journey time is reduced on the existing route while frequency stays at today's eighteen calls, the station stays where it is, and fares stay conventional. Nothing changes except speed.

TABLE 7 · RIDERSHIP AT CONSTANT FREQUENCY, EXISTING STATION

Toronto journey time	Time saved	Annual trips	Change
135 min, as operated today	—	450,000	—
118 min, reliable 160–177 km/h	13%	471,000	+4.6%
95 min, HPPR class	30%	502,000	+11.6%
80 min, upper bound on this station	41%	525,000	+16.6%

Frequency held at eighteen calls a day, existing station, conventional fares. The final row applies the line-haul time assumed for high-speed operation to the existing station; it exists to separate speed from access, frequency and fare, not as a proposal.

The comparison that matters

Eighty minutes to Toronto from the existing station, at eighteen calls a day and conventional fares, returns about +17 per cent. The same eighty minutes from a station twenty-seven minutes out of town, at eight calls a day and a 25 per cent premium, returns about −17 per cent. The line-haul time is identical. The outcomes differ by thirty-four points, and the whole of that gap is access, frequency and fare — the decomposition set out in Table 3.

The return to speed is sublinear

Each tier delivers roughly 0.4 points of ridership for every point of journey time removed: a 30 per cent time cut buys about 12 per cent more trips. For Kingston's dominant segments, in-vehicle time is a minority of generalised cost — fare, access and waiting make up the rest, and speed touches none of them. The gain concentrates where value of time is highest: on the HPPR-class tier, public sector and institutional travel grows about 17 per cent, general business and leisure 13, seniors and leisure 12, students 10.

A conservatism worth naming

The journey-time elasticity implied by these figures is about −0.4, below the −0.6 to −0.9 range usual in the rail literature, because the fare term in generalised cost dampens the response. A conventional −0.75 applied to the same 30 per cent reduction gives about +22 per cent rather than +12. Table 7 should therefore be read as a lower bound, with the HPPR-class tier plausibly running from +10 to +25 per cent. The earlier comparisons are unaffected, being like-for-like, but the upside of upgrading in place may be larger than shown.

Fare erodes the gain quickly

A ten-point fare increase to help fund an upgrade reduces the HPPR-class gain from about +12 per cent to about +7 — two-fifths of the speed benefit consumed by a 10 per cent ticket rise. That is the same mechanism that sinks the high-speed scenarios, operating here on the alternative, and it is an argument for funding an upgrade from the capital side rather than the farebox.

5 · What the results mean

A stop is not the same as service

The sign of the result is set by three things: how far the station is from where people are going, what the ticket costs, and whether today's service survives. Frequency cannot rescue the outcome alone — at parity with today's eighteen calls, an out-of-town station at a premium fare still returns about −9 per cent. A debate about whether Kingston gets a station, and how many trains stop, is a debate about the wrong variables.

The two claims in tension resolve unfavourably either way

A station justified by strong ridership but served by a minority of trains has its service pattern set by the express timetable, not by the demand said to justify it. Table 6 shows why that is not a compromise: the frequency needed to make the station work at a premium fare is well above what the express pattern tolerates. The international answer is a two-tier service, express and semi-fast, requiring overtaking capacity at the intermediate station — and whether the cost estimate includes it is a question with two answers, both informative.

A conventional comparator

HPPR, the new-build railway proposed under the High Performance Rail framework, is designed for a 240 km/h maximum and 200 km/h typical operation. It serves Kingston without relocating the station, without the fare premium high-speed operation requires, and without reducing the number of trains that call. Table 7 quantifies what that is worth: a smaller share of the theoretical time saving, and a larger share of the ridership.

What would change these results

- A station materially closer to the core, or a committed frequent connection to it, timed to the trains.
- Conventional fares on Kingston flows, which Table 6 shows is the decisive variable.

- A binding commitment that Kingston Subdivision service is maintained, converting the substitution state into the coexistence state, and a published service pattern, so that frequency is a fact rather than an input.

Verification required

Four inputs should be replaced with released data before these figures are cited as more than an order of magnitude: station-level boardings and alightings with origin–destination pairs, which would fix both the base volume and the market split; the calling pattern assumed for the Toronto–Ottawa segment; the fare structure for intermediate stations; and the station location, with its assumed access time to the core. The first is a matter for the operator; the rest sit with the project and its joint project office, whose report and business case remain unpublished.

The finding that matters is not that a Kingston station would fail. It is that the service Kingston already has is the benchmark the project has not been asked to beat — and on the assumptions set out here, it does not beat it.