HIGH PERFORMANCE RAIL · CHAPTER 3

The HPR Concept: Untangling the Corridor

Two incompatible traffic types share one corridor. HPR separates them; ALTO's passenger-only line cannot.

Chapter 2 established that ALTO is the wrong instrument for a real need, and that its deficiencies are founding choices rather than implementation failures. This chapter develops the alternative — but it begins further back than the design itself, because the corridor's problem cannot be understood without first understanding the freight railway that owns it. North America runs a model of rail almost opposite to Europe's, and the company that dispatches the Toronto–Montréal corridor is a profit-maximising private incumbent with capacity problems of its own. Only against that backdrop does the organising idea of HPR — separation — become legible as the right move rather than merely a cheaper one. The argument here is conceptual; the numbers that carry it are developed in Chapters 4, 5 and 8.

3.1 Two Models of Freight Rail: Europe and North America

The corridor debate in Canada is usually conducted with European reference points — the TGV and the ICE — because those are the systems the public pictures when it hears "high-speed rail." But the corridor that any Canadian passenger scheme must share or parallel is a North American freight corridor, and North American freight rail operates on a model that is close to the inverse of Europe's. Getting this contrast right is the precondition for getting the corridor right, because it determines who controls the track and whose traffic comes first.

In North America, the major railways are privately owned, vertically integrated — they own the track and run the trains on it — and profitable. Freight is the core business and holds operational priority; intercity passenger service, where it exists, is a tenant on freight-owned rails. The railways have spent four decades investing specifically in freight productivity: progressively longer and heavier trains, higher axle loads, and double-stack container trains that carry two containers on a single well car. The result is one of the most cost-efficient heavy-freight systems in the world, optimised around freight priority on track the operator owns outright.

Europe is structured almost in opposition. Railways are predominantly state-owned or state-supported; infrastructure and train operations are frequently separated under open-access regimes; and the network's investment and priority have gone overwhelmingly to passenger service, conventional and increasingly high-speed. Passenger trains account for roughly 79 percent of train-kilometres on the European network and hold priority over freight, and since 1990 the length of high-speed line has grown several-fold while the total network has slightly contracted. European rail freight is, in this setting, the marginal — and often loss-making or subsidy-dependent — user of a passenger-first system.

Table 3.1 — Two models of freight rail: Europe and North America

	North America	Europe
Ownership & structure	Private, vertically integrated — owns the track and runs the trains	Predominantly state-owned or state-supported; infrastructure often separated from operations under open access
Operational priority	Freight has priority	Passenger has priority (~79% of train-km)
Investment focus	Freight productivity — train length, axle load, clearance	Passenger service, especially high-speed rail
Maximum train length	Routinely 3+ km; longest exceed 4 km	~740 m standard (835–850 m on a few routes)

	North America	Europe
Double-stack containers	Standard; clearances raised network-wide	Largely foreclosed by loading gauge and catenary height
Rail's freight share	~40% of long-distance ton-miles — among the world's highest	~17% of inland tonne-km, and declining
Commercial position	Profitable; freight is the core business	Freight marginal, often subsidy-dependent
Relative productivity	Low operating cost per tonne-mile	Several times more trains for the same tonnage

Sources: Association of American Railroads (2025); Eurostat (2023–24); comparative freight-rail literature. Figures are order-of-magnitude and current as of 2023–25.

The consequence is a large difference in how much freight each system moves by rail. In the United States, rail carries roughly 40 percent of long-distance freight by tonne-miles — more than any other mode and among the highest shares of any advanced economy (Association of American Railroads, 2025). In the European Union, rail's share of inland freight has fallen to about 17 percent by tonne-kilometres (16.9 percent in 2023, down from roughly 19 percent a decade earlier), and to barely 5 percent once dominant short-sea shipping is included (Eurostat, 2023–24). The gap is wide and, if anything, widening: the European rail-freight share has declined steadily over the past decade while the North American share has held near its long-run high. Detailed decompositions of the difference attribute most of it to structural factors — geography, haul distance, and commodity mix — but a material part to policy, principally Europe's prioritisation of passenger service over freight. The productivity contrast is starker still: moving a given tonnage requires several times as many trains in Europe as in North America, at operating costs per tonne-mile that are a multiple of the North American figure.

3.1.1 Train Length and Double-Stack: The Efficiency the Freight Model Bought

The clearest physical expression of the North American freight model is the size of the trains it runs. Over four decades the continent's railways invested steadily in train length — extending siding loops, spacing distributed-power locomotives through the consist, and adopting heavier couplers and cars — so that a single crew now moves far more tonnage per departure. North American freight trains routinely run well beyond three kilometres, and the longest exceed four, against a European standard freight train of roughly 740 metres (extended to 835–850 metres on a handful of upgraded routes). Because crew and locomotive costs are largely fixed per departure, length is leverage: a train several times longer spreads those fixed costs across several times the payload. The European limit is not a preference but a constraint — siding lengths, signalling block design, and a network built around passenger service cap how long a freight train can be.

The second and larger efficiency is double-stack. Introduced in North America in 1984 and now carrying the majority of US intermodal traffic, double-stack operation places two containers vertically in a single well car, roughly doubling the container payload of a train of given length and cutting line-haul cost per container by something on the order of a third to a half. It required sustained capital investment in vertical clearance — raising or lowering through tunnels, overpasses, and bridges across the network to admit the taller loading gauge — but the North American railways made that investment because the freight payoff justified it. Europe, by and large, did not: its more restricted loading gauge, lower structures, and — critically — overhead electrification installed at heights that predate double-stack and were never designed to clear it, together foreclose the technique across most of the network. The contrast is exact. North America raised its clearances for freight; Europe electrified for passengers at heights that lock freight out of its single largest efficiency.

Two implications carry forward. First, these efficiencies are categorically unavailable to a passenger-only line: train length and double-stack are freight techniques, and a dedicated high-speed passenger railway captures none of them. Second, clearance is destiny for corridor freight economics — a point that returns in section 3.2.4, where vertical clearance on CN's Kingston Subdivision is shown to be the gate between today's single-stack operation and the double-stack economics that would transform the corridor's freight competitiveness. The efficiency the North American model bought with clearance investment is precisely the efficiency a separation scheme can extend to this corridor — and a passenger-only scheme cannot.

~40% / ~17% Rail's share of freight — US long-distance ton-miles vs EU inland tonne-km (2023–25); the US share is the highest of any advanced economy, the EU share is declining	**~79%** Share of European train-km that is passenger service — holding priority over freight on a passenger-first network	**4+ km** Longest North American freight trains, against a European standard of ~740 m — train length is the fixed-cost leverage a freight-priority network is built around

This is the fact the ALTO debate has largely set aside. A high-speed passenger line modelled on the European experience implicitly imports the European settlement — passenger priority, public ownership of the line, freight as an afterthought — into a corridor governed by the North American settlement, in which a private freight railway owns the track, prioritises its own tonnage, and answers to shareholders. The two models cannot simply be overlaid on one another. Any credible passenger proposal for this corridor must begin from the freight economics that actually govern it, not from the passenger-first assumptions of a different continent. HPR begins there. ALTO does not.

> *North America and Europe run opposite railways. In Europe the network serves passengers and freight is the marginal tenant; in North America a private, profitable freight railway owns the track and passenger trains wait in its sidings. A passenger plan that ignores which continent it is on will solve the wrong problem at the wrong cost.*

3.2 CN in This Context: The Economics of the Incumbent

The North American model is not an abstraction on the Toronto–Montréal corridor. It has a name: Canadian National Railway. CN owns and dispatches the Kingston Subdivision, the spine of the corridor, and it is close to a textbook case of the North American freight railway — privatised, relentlessly optimised, profitable, and, decisively, constrained. The company that owns the track decides what runs on it and when, so CN's economics are the corridor's economics. The four subsections below set out CN's value-creation history, the reasons its market valuation has stalled, the difficulty it faces adding capacity, and the intermodal economics that determine where rail beats trucking on this corridor. The account draws on the Initiative's CN research note (April 2026).

Figure 3.1 — CN market capitalisation versus operating ratio since privatisation, 1995–2025

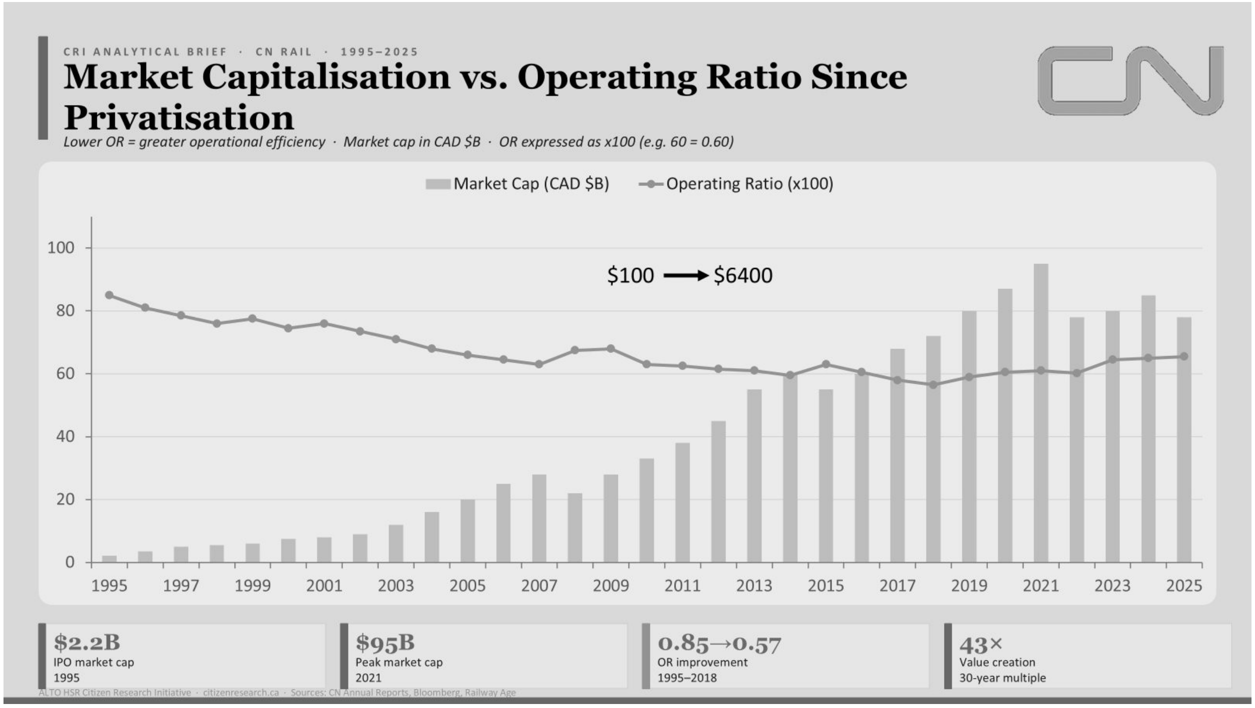

Market capitalisation in CAD billions; operating ratio expressed ×100. The inverse relationship is this chapter's central premise: as the operating ratio fell, market capitalisation compounded — value created by operating discipline on existing track, not by network construction. Sources: CN annual reports; Bloomberg; Railway Age.

3.2.1 Privatisation and the Value-Creation Story

Before the 1990s, CN was a Crown corporation losing money, overstaffed, and without strategic focus. The restructuring begun under Paul Tellier in 1992 and the 1995 initial public offering — then the largest in Canadian history — imposed capital-market discipline for the first time. The 1998 acquisition of Illinois Central reoriented the network toward the north–south trade flows that the 1994 free-trade agreement had unlocked, and brought Hunter Harrison into the company.

The fiscal burden of Crown ownership, 1923–2026

HISTORICAL SUBSIDY (1923–1993)	PROJECTED LOSSES WITHOUT REFORM	COUNTERFACTUAL COST (1995–2026)
$96B	**$1.5B/yr**	**$35–50B**
Total federal support in 1995 dollars — averaging $1.37B per year over 70 years of Crown operation.	CN's own chief executive warned in 1993 that annual losses would reach $1.5B by 1998 without drastic change.	Estimated combined losses, forgone tax revenue, forgone dividends, and capital injections over 31 years.

Sources: Hansard Transport Committee (May 1995); Boardman, Laurin, Moore & Vining (2009); ScienceDirect (2012).

From 2003, as chief executive, Harrison applied Precision Scheduled Railroading at continental scale: moving individual railcars on fixed schedules rather than holding trains for full loads, eliminating classification-yard dwell, shrinking locomotive and car fleets, cutting headcount, and lengthening trains. CN's operating ratio — operating expenses as a percentage of revenue, the defining efficiency metric, where lower is better — fell from roughly 76 percent to about 62 percent in six years. His successor embedded the discipline as institutional practice, and the operating ratio reached a record low of 56.4 percent in 2016, the best sustained margin of any Class I railway to that point. Split- and dividend-adjusted, CN's shares appreciated roughly sixty-fold from the IPO.

The lesson the market drew from that history is the single most important premise of this chapter. CN's extraordinary return was an operational story, not a construction story: value flowed from extracting more output from existing right-of-way, not from building new railway. In rail, the capital-light path to value is efficiency on track you already own. Greenfield construction is the most expensive and most risk-laden route to improvement — a finding that bears directly on the comparison between ALTO and HPR.

> *CN's sixty-fold return since privatisation was earned by squeezing more out of existing track, not by building new line. The North American rail market's clearest lesson is that the capital-light path to value is efficiency on assets already in the ground — the opposite of a greenfield megaproject.*

3.2.2 Why the Valuation Has Stalled

CN's share price has not made durable progress since 2022, and the reasons are structural rather than cyclical. Under a growth-over-margins strategy late in the 2010s, the operating ratio drifted back toward 66 percent, eroding the premium the PSR era had earned. The most consequential setback was the failed attempt to acquire Kansas City Southern: the US Surface Transportation Board rejected CN's proposed voting-trust structure, Canadian Pacific ultimately took KCS, and the resulting CPKC now offers single-line service across Canada, the United States, and Mexico that CN cannot replicate. Network-expansion optionality, one of the two pillars of CN's historical premium, was foreclosed.

A margin-discipline reset from 2022 restored the operating ratio to about 60 percent, with a record quarterly result late in 2025 — but improvement through cost-cutting is familiar to markets and largely priced in, and the easy PSR gains were captured before 2016. Tariff headwinds have been quantified by management at roughly $350 million of revenue with broadly flat volumes, leaving a margin-extraction story rather than a growth story, to which markets historically assign lower multiples. The shares set an all-time high of US$126.64 in March 2024 and, as of mid-2026, trade near US$114 — below that peak and within a 52-week band of roughly US$91 to US$122. The arc below tracks the relationship between operating ratio and share price across the privatisation era.

Table 3.2 — CN operating ratio and share price across the privatisation era

Year	Oper. ratio	Key event	Share (USD)
1995	~87%	IPO — Crown corporation privatised	$2
2003	~70%	Harrison becomes CEO; full PSR rollout	$11
2007	~63%	PSR embedded; best-run Class I railway	$26
2016	~56%	Record-low operating ratio	$72
2021	~66%	KCS bid fails; CPKC emerges as competitor	$112
2024	~63%	All-time high March 2024 ($126.64)	$122
2026	~60%	Margin reset holds; volumes plateau; tariff drag	~$114

Share prices are approximate, split- and dividend-adjusted, in USD (CNI, NYSE). Operating-ratio figures are adjusted where available. Sources: CN annual reports; Macrotrends historical price data; current market data (mid-2026). The inverse correlation between operating ratio and valuation is pronounced through the PSR era; the post-2017 divergence — operating ratio rising while the stock climbed on buybacks — proved unsustainable.

3.2.3 The Capacity Problem and the Milton Precedent

CN's value path is efficiency on existing assets, but growth still requires capacity — and adding capacity by building new infrastructure in populated southern Ontario has proven extraordinarily difficult. CN's three Greater Toronto intermodal terminals run at capacity, and its largest, at Brampton, is hemmed in by surrounding development. The company's answer was a new intermodal terminal at Milton, and its fate is instructive.

CN proposed the Milton Logistics Hub in 2015, having floated and then abandoned a Milton facility between 2001 and 2008. The project drew sustained opposition from the Town of Milton, the Region of Halton, Conservation Halton, and grassroots groups, on grounds of air quality, noise, and the roughly 1,600 daily truck movements a 24-hour terminal would generate. After a federal environmental assessment, the project was approved by Order in Council in January 2021 subject to more than 300 conditions. A Federal Court found the approval unreasonable in March 2024; the Federal Court of Appeal reversed that ruling in October 2024; and the Supreme Court of Canada declined to hear a further appeal in May 2025. Construction is now underway, with operations anticipated around 2027 — roughly a decade later than the originally intended groundbreaking, on a facility costing on the order of a quarter-billion dollars.

THE MILTON LOGISTICS HUB — A DECADE OF FRICTION

Proposed 2015 (after an earlier 2001–2008 attempt was dropped) · federal environmental assessment · approved by Order in Council January 2021 with 300-plus conditions · opposed by the Town of Milton, the Region of Halton, Conservation Halton, and resident groups over air quality, noise, and ~1,600 daily truck trips · Federal Court found the approval unreasonable (March 2024) · Federal Court of Appeal reversed (October 2024) · Supreme Court declined leave to appeal (May 2025) · construction underway, operations anticipated ~2027 — roughly a decade after the intended start, on a ~$250M facility.

The lesson generalises well beyond Milton. Even a profitable, sophisticated incumbent with a clear commercial need cannot readily add freight capacity through new terminal or corridor infrastructure in a populated landscape; community friction is decisive, and it is measured in years and litigation. This is the same friction the Initiative's Community Friction Index captures for the ALTO greenfield corridor in Chapter 2, and it cuts squarely in HPR's favour. The cheapest and least-contested new freight capacity is not capacity built new — it is capacity liberated on track that already exists. That is precisely what separating passenger traffic off the Kingston Subdivision delivers, and it is the foundation of the dual-asset argument developed later in this chapter.

3.2.4 Intermodal Economics, Contrecœur, and the Kingston Sub

Where CN competes for corridor freight is principally intermodal traffic, and intermodal economics are governed less by the speed of the rail line than by drayage — the short-haul truck movement at each end of the journey, from shipper's dock to origin terminal and from destination terminal to consignee. Drayage is a fixed cost that every intermodal move must absorb before rail line-haul savings can generate a net advantage over direct trucking, and in the Ontario–Quebec corridor it imposes an initial disadvantage on the order of $700–900 per container. The crossover distance at which intermodal becomes cheaper than trucking is approximately 820 kilometres on this corridor — longer than the US "500-mile rule" because Canadian drayage labour costs more and the network is less dense.

The decisive implication is that below roughly 500 kilometres, drayage and terminal handling account for 60 to 85 percent of total intermodal cost, and the efficiency of the rail line itself is almost irrelevant. Terminal siting and corridor clearance, not line speed, determine freight competitiveness. Two levers move the economics materially. The first is on-dock rail, which eliminates a drayage leg entirely and shifts the crossover distance from roughly 820 kilometres to roughly 470 — opening a far larger share of corridor freight to rail. The Montréal Port Authority's Contrecœur terminal, about 40 kilometres east of the port and engineered for direct CN rail access at the dock with colocated logistics land and appointment systems from the outset, is the Canadian benchmark for this design. The second lever is double-stack operation, which roughly halves the line-haul cost per container and moves the crossover down a further 150 to 200 kilometres — but requires adequate vertical clearance throughout the corridor. On CN's Kingston Subdivision, clearance is a material constraint and a distinct category of capital investment.

DRAYAGE AND THE CROSSOVER DISTANCE (ONTARIO–QUEBEC, 2025 CAD)

Rail is ~40–50% cheaper than truckload per tonne-km at distance, but drayage is a fixed penalty regardless of haul length. Crossover where intermodal beats trucking: ~820 km under current assumptions (truck ~$2.85/km, rail line-haul ~$1.45/km, drayage ~$380/leg, terminal handling ~$180/container). Below ~500 km, drayage and terminal handling are 60–85% of intermodal cost — line speed barely matters. On-dock rail (one dray leg eliminated) shifts the crossover to ~470 km. Double-stack halves per-container line-haul and moves it down a further 150–200 km, conditional on vertical clearance. Contrecœur (1.15M TEU at full build, ~2030; direct CN rail at dock) is the design benchmark.

These two levers point to the same conclusion. The corridor's freight economics improve on uncongested freight paths and on double-stack clearance — the first reducing the operating friction that drives the operating ratio, the second unlocking the line-haul efficiency that the North American model is built around. Both bear on metrics CN's own public reporting treats as central: the operating ratio it reports quarterly, and the terminal and network capacity constraints documented in section 3.2.3. A scheme that frees Kingston Subdivision paths of passenger traffic and is planned with clearance in mind would therefore act on the same variables CN's published strategy is directed at. CN has taken no public position on HPR, and nothing here should be read as its endorsement; the point is that the incentives run in the same direction. CN is not the only freight railway on this axis, however, and the next section turns to the other before the chapter states the conflict that currently prevents the corridor from working.

3.3 What About CP?

The corridor analysis so far has centred on CN, for a simple reason: CN owns and dispatches the Kingston Subdivision, the line on which the corridor's intercity passenger service runs. But CN is not the only Class I freight railway on the Toronto–Montréal axis. Canadian Pacific — since the 2023 Kansas City Southern merger, Canadian Pacific Kansas City, or CPKC — runs a main line of its own through the corridor. It is reasonable to ask whether CPKC changes the picture. It does not, and the reason is geographic.

CPKC's centre of gravity runs the other way. Its principal Southern Ontario mainline is the Galt Subdivision, west from Toronto toward Windsor and the Detroit gateway, and since the KCS merger its strategic priority has shifted decisively north–south — to the continental Canada–United States–Mexico network that is now the primary growth story in every CPKC investor communication. On the Toronto–Montréal axis CPKC runs its own line, the Belleville Subdivision, which joins the CN Kingston Subdivision near Newcastle, runs roughly parallel to it through the central part of the corridor, and turns north east of Belleville toward Smiths Falls, continuing to Montréal on CPKC's own metals. CPKC's eastern freight, in short, travels on its own track, not on CN's.

This is decisive for the argument that follows. The passenger–freight conflict the next section describes is specific to the CN Kingston Subdivision, where VIA's intercity trains share track with CN's freight; CPKC, running on the separate Belleville Subdivision, is not a party to it, and the freight capacity that separation liberates accrues to CN, the line's owner — the freight dividend this chapter develops later. The Belleville Subdivision matters in one further respect: it is a second existing rail right-of-way running parallel to the CN line and Highway 401 through the corridor's central section, additional evidence that the ground HPR would follow is already a multi-track, multi-modal transport spine rather than open country — the disturbed-corridor condition on which the brownfield case in section 3.5 depends. CPKC's own history, finally, echoes this chapter's central lesson: privately built from 1881, transformed after 2012 by the same Precision Scheduled Railroading discipline that remade CN (its operating ratio cut from above 80 percent to under 60), and made continental by the KCS merger, CPKC created its value through efficiency and integration on assets that already existed — not through new construction.

3.4 The Entanglement Problem

Sections 3.1 to 3.3 establish the actors. This section states the conflict between them. The corridor's mobility failure, documented in Chapter 2, has a single physical cause: intercity passenger trains and heavy freight trains share one set of tracks — principally the CN Kingston Subdivision between Toronto and Montréal — and they are not compatible occupants of it. The on-time performance below 64 percent, the four-hour-plus Toronto–Montréal schedule, the inability to run frequent service: these are not separate problems. They are symptoms of one underlying condition, which is that two traffic types with opposite operating characteristics are competing for the same capacity on infrastructure that belongs to, and is dispatched by, one of them.

Freight and passenger trains want different things from a railway. Freight is heavy, long, and slow to accelerate; it tolerates delay, runs to commercial rather than clock-face schedules, and imposes high axle loads that shape track standards and maintenance cycles. Intercity passenger service is the inverse: light, short, quick to accelerate, intolerant of delay, and valuable precisely in proportion to its frequency and reliability. A track configuration optimised for one is actively hostile to the other. Run them together and each degrades the other's economics.

The arrangement is not symmetrical, and that is the decisive feature — the feature sections 3.1 and 3.2 were written to establish. CN owns the Kingston Subdivision, dispatches it, and runs freight on it as the core of a business it has spent thirty years optimising to extract maximum tonnage at minimum cost. VIA Rail is a tenant. Under the host-railway priority dynamic that governs almost all shared track in North America, when a freight movement and a passenger movement want the same segment at the same time, the freight train proceeds and the passenger train waits in a siding. VIA cannot unilaterally add frequency, because every additional passenger path must be negotiated against the freight schedule of the railway that controls the dispatching. The passenger operator does not hold the levers that determine its own reliability, and under this ownership structure there is no commercial mechanism that would require the track owner to surrender freight capacity to improve them.

Seen this way, the corridor's failure is one problem with two faces. Every passenger path VIA runs is a path CN cannot use for freight, and every siding meet and priority negotiation is friction on a network built and operated to move tonnage without it. The two traffic types are entangled: neither can be optimised without disadvantaging the other, because they draw on a common and finite pool of capacity owned by the party whose traffic comes first. This is the entanglement problem, and it is the problem to which both ALTO and HPR are answers.

> *Passenger trains wait in sidings because a profit-maximising freight railway owns the track and its tonnage comes first. The corridor's failure is not a shortage of speed — it is a structural conflict over capacity that no schedule adjustment can resolve while the two traffic types remain entangled on one set of rails.*

3.5 Separation as the Design Principle

If entanglement is the disease, separation is the cure — and it is the single design principle from which the entire HPR proposal is derived. The aim is to give intercity passenger service its own dedicated path, engineered to passenger requirements, and in doing so to hand the shared corridor back to freight. Each traffic type then runs on infrastructure suited to it, and the capacity conflict that produces the corridor's present failure simply ceases to exist. Crucially, this resolves the conflict in a way consistent with the North American settlement of section 3.1 rather than against it: it does not ask a private freight railway to subordinate its tonnage to passenger priority, as the European model would; it removes the passenger trains from the freight railway's tracks altogether. Everything technical in the chapters that follow — alignment, geometry, cost model, phasing — is an implementation of this one idea.

Separation can be achieved cheaply or expensively, and the difference between the two is the difference between HPR and ALTO. HPR pursues it through a brownfield-led, incremental philosophy: build the dedicated passenger path along the geometry of corridors that are already disturbed — the Highway 401 transportation corridor and the existing rail right-of-way, joined through the corridor's central section by the parallel CPKC Belleville Subdivision noted in section 3.3 — upgrading and reusing infrastructure wherever the engineering permits and building new only where geometry or capacity genuinely demands it. This is not merely a cost-control measure. It is a network-optimisation move: it treats the corridor as a system to be reorganised rather than a blank sheet on which to draw a single new line, and it keeps the new passenger alignment adjacent to the freight corridor it is relieving rather than divorcing the two.

Electrification follows the same logic of fitting the investment to the demonstrated case. HPR electrifies where traffic density and operating economics justify it and stages the transition elsewhere, rather than treating full electrification as a fixed requirement of the specification. The traction question — diesel, electric, and staged electrification as a pathway — is examined in detail in Chapter 4; what matters here is that electrification is treated as a parameter to be optimised, not a banner to be planted.

The scope of the HPR triangle is best understood as two distinct elements rather than a single undifferentiated network, and the report keeps them separate throughout because they carry different costs, risks, and engineering profiles. The first is the spine: a measured corridor of roughly 479 kilometres on the Toronto–Montréal axis, running from Pickering Junction in the east-GTA approach to Dorval at the Montréal end, built and upgraded as a dedicated passenger path. The second is the Ottawa legs: roughly 200 kilometres of upgrade over the VIA-owned Smiths Falls and Alexandria Subdivisions connecting the capital into the network. The spine is where the new-build cost concentrates; the Ottawa legs are an upgrade of track that the public passenger operator already owns, and are costed separately and at a materially lower rate in Chapter 4.

THE TRIANGLE, TWO ELEMENTS

Element one — the spine: ~479 km, Pickering Junction–Dorval, the Toronto–Montréal axis. Dedicated passenger path, brownfield-led, built and upgraded along disturbed-corridor geometry. This is where the HPPR spine cost is grounded (Chapter 4): about $19.0B as specified and $26.1B on the de-biased central estimate — roughly $40M/km as specified, $54M/km de-biased — for the 479 km diesel configuration including full grade separation, with electrification a deferrable ~$3.1B option. Element two — the Ottawa legs: ~200 km over the VIA-owned Smiths Falls and Alexandria Subdivisions. Existing publicly owned track, upgraded rather than newly built, costed separately at a lower rate. The two elements are kept distinct in every cost, ridership, and risk figure in this report.

Separation, pursued this way, is structurally different from ALTO's approach in a respect that the next two sections develop. Because HPR builds along the corridor it is relieving, the act of removing passenger trains from the shared track is also the act of liberating freight capacity on that same track — capacity that, as section 3.2 showed, the corridor's owner has direct commercial reason to want. The passenger investment and the freight benefit are two outputs of one build. That is the dual-asset structure introduced in Chapter 1, and it is the subject of section 3.6.

3.6 Build One, Make One Free

The motif that runs through this report — build one, make one free — has until now been asserted and framed. Chapter 1 stated it in the executive summary and bonded it to the megaproject literature, establishing that a dual-asset structure is a response to the benefit-shortfall failure mode that sinks megaprojects. This section unpacks the four words in full and shows the mechanism that makes them literally true rather than rhetorically convenient.

It helps to give the two assets their names. The proposal as a whole is High Performance Rail, or HPR. It comprises two components, and the four words of the motif map onto them exactly: "build one" is High Performance Passenger Rail (HPPR), and "make one free" is High Performance Freight Rail (HPFR).

Build one. The single thing the project builds is HPPR — a dedicated intercity passenger path along the corridor. That is the capital project: the spine and the Ottawa-leg upgrade described in section 3.5, engineered to passenger requirements, paid for as a passenger investment, and the configuration the capital cost estimate prices. Its direct return is the passenger service — faster, more frequent, more reliable intercity rail on the Toronto–Ottawa–Montréal triangle. HPPR is the asset that appears on the balance sheet and the asset whose benefits the financial model in Chapter 8 must justify against its cost.

Make one free. The moment passenger trains move onto the HPPR path, they vacate the Kingston Subdivision. The capacity they were consuming — the paths, the priority negotiations, the siding meets, the schedule constraints that freight had to work around — reverts to the corridor's owner as usable freight capacity: HPFR, the uncongested and — where clearance is addressed — double-stack-capable freight corridor that the Kingston Subdivision becomes once the passenger conflict is removed. A second asset has been created: not a new railway, but liberated, higher-performance capacity on an existing one. And it has been created at no incremental cost, because no additional construction was required to produce it. It is a byproduct of the geometry of the build. HPPR and HPFR are not two projects that happen to be bundled; they are one project — HPR — with two separable outputs, the second of which falls out of the first for free.

This is why the design principle of section 3.5 matters so much. HPFR is available only because HPPR's dedicated path runs along the corridor it relieves. The separation is what converts a passenger investment into a dual-asset investment: by taking passenger traffic off the shared track, HPR simultaneously solves the passenger problem and the freight problem that section 3.4 showed were the same problem seen from two sides. One build, two assets. Build one, make one free.

> *The freight-capacity asset is not a benefit HPR claims by accounting convenience. It is a physical consequence of moving passenger trains onto their own line: the paths they vacate on the shared corridor become usable freight capacity the instant they leave. One capital project produces two assets, and the second costs nothing to create.*

The connection to the report's analytical discipline, established in Chapter 1, is worth restating because it is the deeper reason the dual-asset structure matters. The benefit-shortfall failure mode — benefits that come in lower than forecast — is one of the two mechanisms by which megaprojects fail. When a single investment generates two separable streams of benefit, the passenger business case need claim only the passenger benefits to clear its threshold; the freight-capacity dividend is a surplus that sits outside the passenger appraisal and is therefore untouched by the reference-class haircut applied to passenger benefits. The dual-asset structure is, in this precise sense, a bias countermeasure: it builds benefit redundancy into the design, so that the case does not depend on the optimistic passenger forecast that the empirical record warns against.

3.7 HPFR: The Freight Dividend

The liberated capacity of section 3.6 is not an abstraction, and — as sections 3.1 and 3.2 established — it lands on a named, motivated beneficiary. It is freight paths returned to a corridor owned by a profit-maximising railway that is actively seeking exactly this kind of value. This section describes who gains and how; the magnitude of the gain, expressed as displaced emissions in the environmental analysis and as a benefit in the financial model, is handed forward to Chapters 5 and 8 respectively.

That dividend is HPFR, the High Performance Freight Rail defined in section 3.6. This is the section where it stops being a structural claim and becomes a description of who gains and how. HPFR is not free-floating capacity: it is value that accrues to identifiable parties, and the most consequential of them is the railway that owns the track. The section takes the beneficiaries in turn — the host railway, then shippers and the wider supply chain, then the public.

The host railway gains first and most directly. CN recovers paths on the Kingston Subdivision previously consumed by passenger service and by the operational friction of accommodating it. Uncongested freight paths translate into higher network fluidity, more predictable transit times, and the ability to grow tonnage without the capacity ceiling that mixed traffic imposes — and, as section 3.2 showed, reduced operating friction bears directly on the operating ratio that is the master variable for rail valuation. Where the separation programme is paired with clearance investment, it also enables the double-stack operation that halves line-haul cost per container. A freight railway that no longer dispatches around passenger priority, on a corridor cleared for double-stack, is a materially more valuable railway on the same physical asset — HPFR in practice.

Shippers and the wider supply chain gain second. Reliable, uncongested rail capacity on the busiest goods corridor in the country is a resilience asset: it raises the ceiling on how much freight can move by rail, reduces the variability that pushes shippers toward more expensive or higher-emission alternatives, and — combined with on-dock terminal design of the Contrecœur type — extends the reach of competitive intermodal service to a larger share of corridor flows. This benefit exists only because the passenger investment freed the capacity that makes it possible.

The public gains third, and this is where the freight dividend connects to the environmental case. Freight capacity that would otherwise be unavailable on rail is capacity that can absorb goods movement currently carried by road. Each tonne-kilometre shifted from truck to rail reduces highway congestion, road wear, collision exposure, and — most consequentially for the report's climate accounting — carbon emissions, given the substantial per-tonne-kilometre advantage of rail over road haulage. The displaced-road-freight emissions benefit is a core component of the HPR lifecycle CO_2 case and is quantified in Chapter 5.

WHERE THE FREIGHT DIVIDEND IS QUANTIFIED

This section establishes the freight dividend qualitatively — who benefits and through what mechanism. Its magnitude is developed elsewhere. The freight-capacity benefit enters the financial appraisal as a distinct benefit stream in Chapter 8 (§8.3), where it stops being rhetorical and becomes a number in the benefit-cost ratio. The displaced-road-freight emissions and the broader CO_2 lifecycle treatment are developed in Chapter 5. The dividend is carried through the report as a measured quantity, not asserted as a talking point.

The essential point for this chapter is structural. The freight dividend is a real, multi-party benefit — to the host railway, to shippers, and to the public — and it exists as a direct consequence of the separation principle. It is the "make one free" half of the motif made concrete, and it falls on a balance sheet already looking for it. A scheme that separates the traffic flows by building along the corridor captures all of it. A scheme that does not is the subject of section 3.8.

3.8 Why ALTO Cannot Do This

ALTO is also a separation scheme — it too gives passenger trains a dedicated line — so it is reasonable to ask why the freight dividend is not equally available to it. The answer turns on what kind of separation ALTO performs, on which continent's model it imports, and at what cost. It explains why the report characterises ALTO as single-purpose spending and HPR as positive-sum.

ALTO's dedicated line is a greenfield passenger-only corridor engineered for 300 km/h and routed away from the existing freight corridor, through new terrain in Eastern Ontario. That choice imports the European passenger-first model — a dedicated high-speed passenger railway as an end in itself — into a North American freight corridor whose economics it does not engage. Three consequences follow, and together they foreclose the freight dividend. First, because the new line does not run along the Kingston Subdivision, ALTO does not reorganise the freight corridor; it builds a parallel passenger facility through different ground and leaves the freight network's configuration, congestion, and clearance constraints essentially as it found them. Second, ALTO is a single-purpose asset: a passenger-only line at 300 km/h cannot carry freight and is not designed to, so the only benefit it can capture is the passenger benefit. Third, and decisively, ALTO's own business case books no freight benefit at all. The freight dividend does not appear in ALTO's appraisal because ALTO's design does not produce it as a captured, usable, or claimed output.

The cost dimension sharpens the contrast. ALTO solves the single problem it addresses — passenger throughput — at the maximum per-kilometre cost the corridor admits: a new greenfield high-speed alignment through sensitive terrain. The capital cost analysis in Chapter 4 places HPPR at roughly $40 million per kilometre as specified and about $54 million per kilometre on the de-biased central estimate — some $19.0B to $26.1B for the 479 km spine — against ALTO at approximately $142 million per kilometre on the same reference-class basis. On a like-for-like de-biased comparison ALTO therefore costs roughly two and a half times as much per kilometre to produce one benefit stream, and considerably more in total once its far longer corridor is counted, while HPR's lower-cost, corridor-aligned build produces two benefits — HPPR and HPFR — and leaves CN's capacity and clearance problems measurably better than it found them. This is the meaning of single-purpose versus positive-sum spend: not a slogan but an accounting of how many separable benefits each dollar of capital buys.

> *ALTO builds one and makes nothing free. A greenfield passenger-only line on the European model solves passenger throughput at the corridor's maximum per-kilometre cost and leaves the North American freight conflict exactly where it found it. HPR builds one and makes one free at roughly a third of the per-kilometre cost. The difference is not in execution — it is in the concept, and in which continent's railway it was designed for.*

This is the structural superiority the chapter set out to establish, and it rests entirely on the design principle of section 3.5. Because HPR separates the traffic flows by building along the corridor it relieves, separation itself becomes the source of a second asset that the corridor's owner has commercial reason to want. Because ALTO separates them by building a remote, single-purpose, maximum-cost line modelled on a different continent's railway, separation produces nothing beyond the passenger service it was built for. Both schemes answer the entanglement problem of section 3.4; only one of them answers it in a way that returns more benefit than it was designed to seek. The remainder of the report demonstrates, in quantitative terms, what that conceptual advantage is worth: the route and engineering that make the corridor-aligned build feasible and the capital cost the brownfield-led approach delivers (Chapter 4), the environmental advantage of a disturbed-corridor route (Chapter 5), the ridership the service can expect on the outside view (Chapter 6), the annual operating economics that ridership supports (Chapter 7), and the financial case in which the freight dividend becomes a number (Chapter 8).

CHAPTER 3 — KEY FINDINGS

The concept in eight parts

3.1 — NORTH AMERICA RUNS THE OPPOSITE RAILWAY

North American freight railways are private, profitable, vertically integrated, and freight-priority; European railways are state-supported and passenger-priority. Rail's freight share is far higher in North America (~40% of US long-distance ton-miles vs ~17% of EU inland tonne-km, and the EU share is falling). A passenger plan modelled on Europe imports passenger-first assumptions into a freight-first corridor.

3.2 — CN IS THE INCUMBENT, AND IT IS CONSTRAINED

CN owns and dispatches the Kingston Sub. Its ~60x return came from PSR efficiency on existing track, not construction — the capital-light path to rail value. Its valuation has stalled (below the March 2024 peak), it struggles to add capacity (Milton took a decade from proposal to construction), and its corridor economics improve on uncongested paths and double-stack clearance.

3.3 — CP IS NOT A PARTY TO THE CORRIDOR CONFLICT

Canadian Pacific Kansas City (CPKC) runs west on the Galt Subdivision and, since the 2023 KCS merger, north–south to Mexico. On the Toronto–Montréal axis it uses its own Belleville Subdivision — parallel to the CN Kingston Sub, then north at Belleville toward Smiths Falls — so it shares no track with VIA. The entanglement, and the freight dividend, are CN's. CPKC's parallel line is also a second disturbed right-of-way in the corridor, and its PSR-and-merger value story reinforces the thesis: value from existing assets, not construction.

3.4 — ENTANGLEMENT IS THE ROOT CAUSE

Passenger and freight share the Kingston Sub and have opposite operating characteristics. Under host-railway priority, freight proceeds and VIA waits. The corridor's failure — sub-64% on-time performance, four-hour-plus schedules, low frequency — is one capacity conflict seen from two sides, on track owned by the party whose tonnage comes first.

3.5 — SEPARATION IS THE DESIGN PRINCIPLE

Give passenger service its own dedicated path and hand the shared corridor back to freight — resolving the conflict with the North American model, not against it. HPR does this brownfield-led, along the 401 and existing rail geometry: a ~479 km Pickering Junction–Dorval spine plus ~200 km of upgrade on the VIA-owned Smiths Falls and Alexandria Subdivisions.

3.6 — BUILD ONE, MAKE ONE FREE

One capital project — HPR — produces two separable assets: HPPR (High Performance Passenger Rail), the dedicated passenger path that is the build, and HPFR (High Performance Freight Rail), the freight capacity liberated on the shared corridor the moment passenger trains vacate it, created at no incremental cost by the geometry of the build. The dual-asset structure is a deliberate bias countermeasure against benefit shortfall.

3.7 — HPFR LANDS ON A MOTIVATED BENEFICIARY

HPFR — the liberated, higher-performance freight corridor — benefits the host railway (fluidity, operating-ratio gains, double-stack where cleared), shippers (resilience, Contrecœur-style reach), and the public (road freight diverted to rail, with emissions and safety gains). It is carried as a measured quantity — a Chapter 8 BCR stream and Chapter 5 emissions — and it falls on a balance sheet already seeking it.

3.8 — ALTO CANNOT CAPTURE IT

ALTO imports the European passenger-only model into a North American freight corridor at ~$142M/km against HPPR's ~$54M/km de-biased (~$40M/km as specified) — roughly two and a half times the unit cost. It does not reorganise the freight corridor, cannot carry freight, and books no freight benefit. ALTO builds one at maximum cost and makes nothing free. The difference between the schemes is conceptual, not a matter of execution.