



ICM
HOLDINGS INTERNATIONAL

RESEARCH & DECISION STUDIO

Corridors are systems

Resolving the chain bottleneck produces more useful capacity than expanding the most visible asset.

Industrial logistics and network capacity

METHODS / EVIDENCE / WORKED STUDIES

ENGLISH EDITION / 2026-10-10

THE DECISION

Corridors are systems

A corridor should be judged by the service it delivers across custody, operating and institutional boundaries. This study derives the serial bottleneck result under explicit annual assumptions, then tests the circumstances in which queues, buffers, batching, calendars, access rights, quality and demand make that result incomplete. Frozen fictional studies reproduce the Corridor lab engine and distinguish a visible asset expansion from relief of the binding system constraint. The analysis translates those contrasts into a measurement dictionary, intervention order, operating envelopes, governance charter and whole-chain commissioning records. Public institutional guidance supplies methods and narrow source claims, not empirical rates for an unnamed project. The conclusion is conditional: prioritise the constraint that governs the declared service, but test feasibility, customer demand and noncompensable obligations before treating greater physical throughput as a better decision.

Questions this study answers

- Which end-to-end service and operating basis must an expansion actually improve?
- What evidence distinguishes an intrinsic bottleneck from starvation, blocking or documentary delay?
- When does a nominal expansion leave delivered throughput unchanged or move the constraint elsewhere?
- Which operating, access, safety, commercial and governance conditions must hold before an intervention can be accepted?

Scope and limits

Illustrative inputs explain mechanisms and are not actual project data or empirical estimates.

A deterministic model cannot establish implementation feasibility, legal effect, lender acceptance or future outcomes.

Cited public guidance is used only within explicitly stated source scope; institutional frameworks require local adaptation.

CONTENTS

Reading the study

01 A corridor delivers a service across boundaries ..	4
02 Capacity claims need an operating basis	10
03 The bottleneck is a system property	16
04 A measurement plan can find the wrong constraint	22
05 Queues convert variability into delay	28
06 Buffers have location and recovery duties	33
07 Batching changes the chain's rhythm	40
08 A port expansion can move the bottleneck inland	45
09 Rail paths differ from rolling-stock capacity ...	50
10 Road capacity is not a truck count	56
11 Maintenance introduces calendar dependence ..	61
12 Weather and seasonality require operating envelopes	66
13 Border processes are part of the production system	71

14 Performance indicators can conceal delivered service	78
15 Commercial incentives can create queues	83
16 Redundancy is useful only if it is usable	89
17 Expansion should follow a bottleneck intervention order	94
18 Governance must cross asset owners	100
19 A commissioning plan must test the whole chain	106
20 When the bottleneck thesis is incomplete	113
21 Appendix: capacity and flow dictionary	119
22 Appendix: illustrative corridor studies	124
23 Appendix: measurement and intervention records	132
24 Appendix: evidence scope and a field mechanism challenge	138

DECISION ARCHITECTURE

01

A corridor delivers a service across boundaries

An end-to-end scope and service definition

Begin with the delivery promise

A corridor investment should begin with a precise statement of the service that its users will receive. A line drawn between a production area and a port is a geography. A schedule of locomotives, yards and loading equipment is an asset inventory. A promise to deliver eligible cargo at an agreed custody point within an agreed operating window is a service. These descriptions overlap, but they answer different questions. Confusing them allows an investment committee to approve a technically impressive component without establishing whether the intended customer can use the resulting chain.

Consider an expressly fictional bulk export corridor. The producer releases cargo, a railway moves it to a coastal stockyard, and a loading operation transfers it to a vessel. The committee sees a loading machine operating intensely during vessel calls and concludes that this visible asset needs expansion. The service manager instead asks whether cargo arrives in the correct grade, with cleared documentation and a compatible vessel ready to receive it. If the railway cannot replenish the yard quickly enough, a faster loader may spend more time waiting. If vessel access is unavailable, inland production may create stock without delivering accepted cargo.

The relevant output therefore needs both a physical and an acceptance definition. Physical movement records the quantity crossing a boundary. Accepted delivery records whether that movement meets the customer's specification and the terms governing custody. Cargo transferred with an unresolved quality dispute is not interchangeable with cargo accepted without reservation. Similarly, a departure from the mine is an upstream event rather than proof of vessel delivery. A throughput report should state which boundary it measures and whether rejected, reworked or diverted material enters its numerator.


This distinction changes accountability. An equipment owner can reasonably warrant performance within a defined envelope without controlling the whole service. A corridor coordinator may be responsible for resolving incompatible schedules while lacking authority to waive

customs or safety requirements. The sponsor should preserve these distinctions in the project design. A broad ambition such as improving logistics does not allocate control, establish a measurement boundary or resolve a conflict between productive movement and lawful release. The opening service statement must make those unresolved matters visible before alternatives are compared.

Exhibit: service boundaries in parallel

Read the relationships as declared service conditions, with the limits stated in the caption.

THE SERVICE HAS PHYSICAL, PERMISSION AND CONTROL TRACKS



A completed local handoff does not establish destination acceptance.

Original conceptual diagram. Solid arrows follow cargo; dotted connections identify eligibility and control. The drawing represents the declared fictional service, with no quantities or measured operating events. Its decision purpose is to expose a completed local handoff that still lacks the condition for destination acceptance.

Scope follows the cargo and its permissions

A useful corridor boundary follows the cargo from its first eligible release to the chosen destination, including transformations that can interrupt acceptance. It also follows the permissions needed for that movement. A train needs infrastructure access as well as rolling stock. A shipment needs documentary release as well as a berth. A road movement needs a compatible vehicle and driver arrangement as well as a nominal road capacity. These dependencies belong inside the service analysis even when the sponsor neither owns nor operates them.

The boundary need not include every remote supplier or institution. Scope should expand when a dependency can change the delivery decision, the operating basis or the ability to recover from disruption. A fuel supplier may belong in a continuity analysis because unavailable fuel can stop the chain. The supplier's unrelated business activities do not. A customer's downstream processing plant belongs inside the delivery

boundary only if its receipt capacity or specification determines what the corridor can acceptably deliver. This decision-based boundary is more useful than either an indiscriminate map or an ownership-only diagram.

An explicit exclusion is an analytical commitment. If the study ends at vessel loading, it cannot describe cargo arrival at the final customer as a demonstrated result. If it omits maritime sailing, the estimated transit time is not a door-to-door time. If the operating scenario excludes a border closure, its apparent capacity cannot establish continuity during that closure. Exclusions should be attached to the conclusions they limit. A distant disclaimer does little to prevent a diagram or executive summary from suggesting a wider claim than its actual boundary supports.

Interfaces deserve particular attention because each side may define completion differently. A mine may record release when cargo leaves a stockpile, while the railway records acceptance when a train is formed and weighed. The interval contains loading, inspection and queuing that neither local report may own. At the coastal end, a yard may count receipt before sampling is complete while the loader accepts only material released for shipment. The system ledger must expose these intermediate states. Otherwise apparent losses or delays become arguments about whose report is correct instead of evidence about the service.

Draw the service before selecting the asset

The service map should show physical progression, documentary progression and control progression. Physical progression identifies where material is held, transferred or transformed. Documentary progression identifies the evidence that permits the next movement. Control progression identifies who can authorise, stop or restore that movement. Keeping these tracks distinct makes it possible to locate a disagreement. A train may be physically ready, administratively ineligible and operationally outside its allocated path at the same time. A single green status hides the reason it cannot depart.

A stockpile is particularly revealing. It may protect the loader from a short rail interruption, yet it may also conceal unreliable inland replenishment. It can contain cargo of several grades, cargo awaiting release and cargo reserved for another vessel. Treating the entire pile as available supply may overstate the service buffer. The initial system map should therefore describe eligibility transitions rather than only storage volumes. This is the beginning of a working definition of useful capacity: the capacity to advance eligible material through the relevant service boundary under the declared conditions.

The map should also identify the user affected by delay. A producer may care about evacuation from the mine, a charterer about vessel readiness, and a downstream buyer about received quality. Their objectives can conflict. Faster evacuation can overload the yard; a vessel priority rule can delay another cargo; tighter acceptance tests can slow release while reducing rejection risk. The committee needs an explicit service hierarchy or a process for resolving such conflicts. An average throughput objective alone cannot decide which user's commitment should prevail during a constrained operating period.

The fictional studies in this paper deliberately simplify this map to annual serial flow. They show the effect of changing annual effective capacities while holding the defined chain and operating cost assumptions fixed. They do not assign custody, simulate individual consignments or prove acceptance. The broader service map supplies the questions that the annual calculation leaves open. In practice, the sponsor should approve the map and service definition before accepting a capacity model, because the model's minimum operation can be mathematically correct for a boundary that does not answer the customer's question.

A disputed delivery exposes the boundary

Imagine a fictional disagreement after cargo has reached the port yard. The producer's report says its evacuation commitment was met. The carrier's report says the transport movement was complete. The yard's receipt record confirms physical arrival, but part of the cargo remains on hold pending an eligibility decision. The vessel-loading actor says the intended parcel is not ready. None of these statements has to be false. They refer to different events, and only the last unresolved transition determines whether the chosen vessel service can proceed.

The coordinator should not resolve the disagreement by selecting the most senior actor's tonnage report. It should reconstruct the agreed service object. Which commercial parcel was intended? Which physical receipt objects compose it? What acceptance condition remains unresolved? Who can provide the missing evidence, and what other eligible work can proceed while it is resolved? The resulting record separates completed local duties from the unmet system commitment. This can support fair responsibility allocation without pretending that the completed upstream movement has no value.

The investment implication depends on that reconstruction. A larger receipt system will not release held cargo if the missing condition is documentary or quality-related. A faster loader will not make the parcel eligible. A change in release preparation may improve the service, while

additional segregation space may protect other cargo from the hold. The case directs the team toward a specific mechanism and feasible alternatives. It also prevents a service failure from becoming a general narrative that every asset is too small.

The committee should retain the disagreement as a test of the service definition. If the parties cannot identify the event that counts as successful delivery, the project's objective remains unresolved. An investment approval should then require a clarified acceptance boundary and governing record before it claims a capacity benefit. This is a practical use of the systems perspective: it makes an ambiguous operational dispute answerable through objects, events, evidence and authority instead of competing local totals.

Define an intervention that can be falsified

The claim that relieving a bottleneck improves the corridor becomes testable only after the target service and comparison basis are fixed. An intervention record should state what changes, what remains constant, which delivered outcome is expected to respond and what observation would contradict that expectation. For a loader expansion, the expectation might be a reduction in time needed to transfer released cargo during compatible vessel windows. The whole-chain throughput claim requires a further showing that another stage, customer demand or a permission does not absorb that apparent gain.

A good falsifier is specific enough to alter a decision. If eligible cargo delivery remains unchanged after a controlled operating trial while inland queues grow, the intervention has not demonstrated the promised service improvement. That observation does not automatically prove that the equipment is defective. It may show that the wrong boundary was selected or that dispatch and stockyard eligibility constrain the benefit. The response should be to revise the system explanation and investment sequence rather than redefine success as local utilisation after the fact.

The service definition should remain versioned. A sponsor may later add another product, destination or operating season. Such a change can be commercially sensible, but it alters the basis on which the earlier capacity conclusion was made. The new service should trigger a reassessment of bottlenecks and operating constraints. Otherwise a model originally intended for a homogeneous flow may be carried into a mixed operation where handling, quality segregation and scheduling matter. A service change is not merely a new label on an unchanged throughput number.

The most useful opening decision is therefore modest and binding: approve the service boundary, acceptance basis, accountable actors and test of improvement. This gives the technical team a coherent question and gives the committee a means of challenging an attractive asset proposal. It also prevents governance from arriving only after procurement, when incompatible obligations are harder to alter. The first deliverable of corridor development is a shared description of what must be delivered and under which conditions, not a larger number attached to the most visible piece of infrastructure.

Exhibit: the same cargo has different states

Original conceptual map for the fictional bulk corridor. Each row describes a boundary that needs an event record; no actual operating data are represented.

Exhibit: the same cargo has different states			
Boundary	Physical state	Permission or acceptance	Accountable decision
Producer release	Cargo ready for collection	Grade and release evidence complete	Producer authorises eligible batch
Rail acceptance	Cargo loaded and consist ready	Access and train formation requirements satisfied	Rail operator accepts movement
Yard receipt	Cargo physically received	Sampling, segregation and shipment release resolved	Yard allocates usable stock
Vessel transfer	Cargo crosses loading boundary	Compatible vessel and cargo acceptance conditions met	Marine and cargo actors confirm transfer

Service-definition gate

Approve the service statement only when the customer boundary and the evidence of successful delivery agree. Record the purpose of the corridor, the eligible cargo, custody transitions, operating envelope and authority to stop. Require each asset proposal to explain which part of this service it changes. Return the proposal for revision when its benefit is stated at a local equipment boundary while the approved objective sits elsewhere.

1. Name the origin, destination and acceptance event.
2. Describe permissions and transformations between those events.
3. Agree which unmet condition prevents the service from being counted.
4. Record who changes the service definition and who reviews the resulting capacity basis.

An ownership boundary is easier to manage

An owner may argue that a wider service boundary creates responsibility for events it cannot control. That concern is valid for contracts and warranties. It does not justify excluding those events from the analytical chain. The map should distinguish influence from authority and should avoid assigning an owner obligations it has not accepted. A sponsor can recognise a border agency as a decisive dependency without pretending to control it. The resulting investment decision may require a condition, an alternate service promise or a governance arrangement rather than a wider equipment guarantee.

A service map establishes a question, not feasibility

The map cannot establish that a proposed route is permitted, that cargo quality can be maintained or that each actor will accept the proposed allocation of responsibility. Those require relevant evidence and agreement. Its role is to expose dependencies and prevent a narrow capacity result from being interpreted as demonstrated end-to-end delivery.

Evidence and scope

The World Bank toolkit describes corridors through connected infrastructure, services and coordinating institutions. Its primer includes gateways within the corridor and argues for attention to total movement from origin to destination. This supports a broad service boundary, while its examples supply no operating values for the fictional chain.

Evidence: Trade and Transport Corridor Management Toolkit

Decision implications

- Define successful delivery before sizing the assets.
- Include decisive permissions and interfaces even when they sit outside the sponsor's ownership.
- Attach each capacity claim to the boundary it actually measures.

DECISION ARCHITECTURE

02 Capacity claims need an operating basis

A units and operating-calendar worksheet

A rate needs a denominator and an envelope

A statement of annual capacity conceals a calendar, a cargo specification and an operating method. The capacity of a loader may originate as a rate under a defined material condition while a railway capacity may originate as a timetable with particular train lengths and turnaround assumptions.

A stockyard figure may describe storage rather than annual flow. These quantities cannot enter the same minimum calculation until the analyst establishes compatible units and operating bases. An impressive nominal number is not enough to make a stage comparable with its neighbours.

The first task is to distinguish a rate from a stock. A rate describes movement per unit of time. A stock describes material held at an instant. Stockyard area or usable storage can influence continuity and batch formation, but neither directly gives an annual throughput without a residence-time and handling basis. Similarly, the number of wagons is a fleet stock, while delivered cargo depends on payload, cycles and availability. Converting one into the other requires assumptions about operation, not a change of unit label.

Nominal capacity should identify the conditions under which it is attainable. These may include the handling configuration, allowable material characteristics, loading sequence, personnel, supply readiness and scheduled hours. A nominal annual quantity that already includes downtime should not be multiplied by an availability factor that removes the same downtime again. The analyst should ask whether the supplier's figure is instantaneous, shift-based or annualised, and whether it assumes uninterrupted feed. That question often matters more than another decimal place in the final throughput result.

The operating calendar should be shared only where it is genuinely common. A railway can be closed for planned maintenance while a port remains open; a vessel-loading operation may depend on tide or weather windows that do not constrain inland movement at the same time. Expressing each stage as an annual rate is a useful simplification, but it does not prove that those opportunities line up. The annual model assumes enough coordination and buffering to make the declared effective capacities mutually usable. Later chapters test the implications of that assumption.

Availability is a definition before it is a percentage

Availability can refer to mechanical readiness, scheduled operating readiness or demonstrated service capability. An asset may be mechanically ready while its crew, fuel, path or cargo is unavailable. A report that counts that period as available can be correct for maintenance management and misleading for delivered service. Before using availability in a model, the team should define the denominator, state exclusions and distinguish the reason an apparently ready asset did not perform. An availability percentage without that dictionary is an argument waiting to happen.

The Corridor lab applies each user-supplied availability to nominal annual capacity. This is an explicit effective-capacity convention, not a reliability model. The resulting figure should be read as an annual usable capacity assumption for the stated stage. It does not assign random failure probabilities or simulate when unavailable periods occur. Two stages can have the same annual availability yet produce very different corridor service if their interruptions are coordinated, alternating or concentrated within the same critical season. The calculation deliberately leaves that chronology unresolved.

An analyst should maintain a downtime reconciliation. Planned maintenance, unplanned repair, missing crew, denied access, cargo starvation and downstream blocking are different mechanisms. Combining them into a single capacity multiplier may be adequate for a screening study, provided the team states what is inside it. It becomes dangerous when an intervention targets only one mechanism but the model assumes it improves the entire multiplier. A mechanical upgrade cannot automatically create a train path or release a consignment held for documentation.

The distinction also prevents circular measurement. If availability is inferred from realised throughput divided by nominal capacity, and realised throughput was limited by demand, the factor measures use rather than capability. Multiplying nominal capacity by that inferred factor simply reproduces historical flow. It cannot tell the committee what the stage could deliver under additional demand. A defensible operating basis separates intrinsic readiness from the system's choice or inability to utilise the stage, then states which concept the effective-capacity input represents.

Cost units must follow the delivered flow

The unit-cost calculation is equally sensitive to basis. Variable cost per tonne should describe the marginal operating cost attached to cargo moving through the stated stage. Annual fixed cost should describe the obligations that remain under the assumed operating regime. Their sum across stages is allocated to the common delivered throughput in the canonical model. This convention makes the denominator effect visible: a chain that delivers less cargo spreads the same assumed fixed obligations over less useful output.

That common-flow convention assumes no material loss, no rework loop and no difference between stage quantity and accepted destination quantity. If a material is dried, blended, rejected or recirculated, local tonnes and delivered tonnes differ. Then stage cost needs a conversion

basis before it can be added consistently. Moisture also makes an apparently simple tonne ambiguous. The team should specify whether a physical measurement is wet, dry or defined by another acceptance convention, and should avoid treating a conversion factor as universally known.

The frozen intervention cases intentionally hold operating costs constant. This isolates what the annual serial model says about capacity and fixed-cost dilution. It does not imply that an expanded railway or loader can be acquired without capital expenditure, incremental maintenance or additional labour. Reading the lower illustrative unit cost as an investment return would skip the economic question. A proper appraisal would require incremental cost, timing, financing and the value of additional accepted service, all of which lie beyond the current app's method.

Currency and price basis need the same discipline. A real estimate should state whether costs use a constant or changing price basis, which currency applies and whether taxes, financing, insurance and replacement obligations enter the calculation. The fictional model's currency labels simply keep the calculation dimensionally coherent. They do not establish tariffs, exchange rates, market costs or contract entitlements. The absence of project-specific commercial information is deliberate, and the conclusions should remain physical screening conclusions rather than disguised transaction advice.

Build a basis sheet that another operator can challenge

The basis sheet should permit a second operator to reconstruct each input without guessing how the first analyst interpreted a label. For every stage it should record the physical boundary, cargo basis, nominal-capacity source, calendar, availability definition and cost coverage. It should identify unresolved assumptions by name. An empty source cell must remain empty or be marked unknown; it should not be replaced with an authoritative-looking estimate solely to keep a model running. A screening estimate can be useful if its assumption status and challenge route are explicit.

The sheet should separate measured, stipulated and illustrative values. A measured input needs a defined observation universe and event dictionary. A stipulated input may reflect a design requirement or operating agreement that has not yet been demonstrated. An illustrative input is chosen to test reasoning and represents no actual operation. These

statuses should survive export, review and later reuse. A copied spreadsheet or scenario file can otherwise lose its fictional context and become an apparently factual baseline in a subsequent committee paper.

Changes to the operating basis should trigger a new comparison rather than a silent edit. If an annual capacity is revised because the operating calendar changed, the team should explain which periods became usable and how the adjacent stages respond. If the change represents a corrected unit error, the earlier conclusion should be withdrawn or superseded visibly. Version control is valuable because it preserves the causal explanation of a result, not because a longer change log creates an appearance of diligence.

The resulting decision is to admit only inputs that are sufficiently defined for the current purpose. A rough estimate can pass a concept-screening gate and fail a procurement or service-acceptance gate. The team should name that difference instead of forcing every stage into a universal certainty standard. What cannot be accepted at any stage is a rate with an unknown denominator or a cost with an unknown quantity basis. Those defects prevent the calculation from answering even the simplest corridor question reliably.

Effective capacity in the canonical model

```
C_effective(stage) = C_nominal(stage) × availability_fraction(stage).
Throughput = min(demand, all C_effective). The multiplication removes only
the unavailability represented by the input definition; it does not
independently resolve calendars, queues or access.
```

Illustrative calculation. Assumptions and sensitivity are published with the model.

Completed illustrative basis record

This record describes the frozen base study and marks operational evidence as absent. It is usable as an input review record without implying that the fictional rates have been measured.

Completed illustrative basis record			
Field	Declared basis	Evidence status	Review consequence
Cargo quantity	Uniform illustrative accepted tonnes at all stages	Assumption; no moisture or loss conversion	Reject reuse for a transformed or rejected cargo flow
Annual demand	20.00 Mt/year	Chosen fictional scenario	Test changed-demand fixture before expansion
Rail nominal rate	14.00 Mt/year	Chosen fictional annual capacity	Request calendar and path derivation for any actual application
Rail availability	85.00 percent	Chosen effective-capacity multiplier	Do not interpret as failure probability
Rail fixed cost	42.00 USDm/year	Chosen operating obligation	Incremental expansion investment is excluded

Operating-basis admission test

Admit a capacity input only after its physical quantity, time basis and readiness definition are stated. Ask whether another reduction already appears inside the nominal figure. Reconcile annual fixed and unit variable costs with the same delivered quantity. Preserve an explicit unknown where an input lacks evidence, then decide whether the planned decision can tolerate that uncertainty.

1. Identify the stage boundary and cargo unit.
2. Write the nominal-rate calendar and exclusions.
3. Separate intrinsic downtime from system starvation and blocking.
4. Label the input as measured, stipulated or illustrative and record the next verification action.

A screening model cannot carry an operational dossier

A concept study does not need the evidential depth of a service warranty. It still needs coherent definitions. The basis sheet can be short and can openly use assumptions. Its purpose is to reveal what would change the conclusion, not to bury an early idea beneath documentation. Where the uncertain calendar determines which stage is binding, however, a more detailed basis is necessary before the committee treats the ranking as informative. The level of effort should follow the sensitivity of the decision rather than the prestige of the asset.

An annual capacity convention cannot settle chronology

Even perfectly documented annual capacities can be mutually incompatible in time. The model has no within-year calendar and no inventory state. It cannot demonstrate that an annual rate will be delivered during a particular season or customer window. That requires an additional time-based operating analysis with declared evidence.

Evidence and scope

The AQUA Book provides a framework for quality assurance of decision-relevant analysis and distinguishes verification from validation. This paper uses that distinction to separate correct execution of the annual formula from the suitability of its operating basis; the guidance does not certify the fictional model.

Evidence: The AQUA Book

Decision implications

- Compare rates only after quantity and calendar bases agree.
- Treat availability as a defined input rather than an unexplained performance percentage.
- Keep physical screening distinct from an incremental investment appraisal.

DECISION ARCHITECTURE

03 The bottleneck is a system property

A capacity constraint map with labelled exceptions

Derive the annual serial constraint

The minimum-capacity result follows from conservation in a restricted serial chain. Every accepted tonne must pass through each stage, and no stage can advance more than its effective annual capacity. Delivered flow therefore cannot exceed any stage capacity. It also cannot exceed eligible

demand under the model's definition. The smallest of those bounds limits the chain. This is an upper-bound argument first; treating the bound as attainable requires further assumptions about coordination, continuity and the ability to hold material between stages.

For the canonical app, those further assumptions are built into the steady-state convention. The chain has a single route, consistent cargo quantity, no processing loss and no time-resolved queue or inventory. Each stage is represented by an annual effective capacity rather than a dispatch schedule. Under that convention, the minimum becomes the calculated deliverable throughput. The term deliverable should retain its illustrative model basis. It does not mean an operator has demonstrated the physical, administrative or contractual ability to maintain that service in an actual corridor.

A binding stage is one whose effective capacity equals the calculated chain flow within the engine's comparison convention. When demand is lower than all stage capacities, the model has no capacity bottleneck. When several stages share the smallest capacity, they are jointly binding. These are different states and should not be collapsed into a single recommendation to improve the weakest asset. Demand limitation points toward a commercial or service question. Tied capacities point toward a coordinated intervention question. A unique constraint supports a single-stage sensitivity test but does not establish its economic desirability.

The derivation explains why local spare capacity is not necessarily waste. A nonbinding stage can provide recovery capability, accommodate batch arrivals or protect a service during another stage's interruption. The annual formula only records the difference between effective stage capacity and common flow. It does not value resilience. Calling every positive difference stranded capacity would therefore overstate the model. The useful distinction is between capacity that cannot raise steady-state flow on its own and capacity that has no justified role under the full service objective.

Use the fictional base case to test direction

Illustrative calculation. Assumptions and sensitivity are published with the model.

The frozen base chain delivers 11.90 Mt/year against illustrative annual demand of 20.00 Mt/year. The binding stage is Rail corridor. Its nominal capacity and availability are explicit assumptions rather than measurements. The result is a deterministic gap under those assumptions. It does not estimate the likelihood of missed delivery, the frequency of congestion or the level of profitable customer demand.

Illustrative. The vessel-loading intervention changes only its nominal capacity. Calculated flow remains 11.90 Mt/year and the binding stage remains Rail corridor. This is the clearest form of the paper's thesis: enlarging a visible stage does not increase annual flow when another unchanged stage sets a lower bound. The comparison is useful precisely because it holds demand, the route and other capacities constant. It says nothing about whether the loader expansion improves berth occupation, safety or recovery during a concentrated vessel window.

Illustrative. The rail intervention produces 12.00 Mt/year and shifts the binding stage to Vessel loading. Relieving the original constraint therefore changes the result, but it exposes a nearby limit. The coordinated rail and loading intervention produces 14.40 Mt/year with Port stockyard binding. This sequence shows why an expansion plan should identify the next constraint before committing to a local improvement. A solution that merely shifts the problem can still be useful, but its extent must be stated honestly.

The operating unit costs change with the common throughput denominator while the input operating costs remain fixed. That isolates a mechanical effect rather than a procurement conclusion. A committee should not infer that the coordinated intervention is the best investment solely because its illustrative operating cost per delivered tonne is lower. The cases omit intervention capital costs, timing and feasibility. They establish the physical ranking under the model's restricted assumptions, then hand the economic and operational questions back to the relevant decision process.

A constraint map needs a cause classification

A measured shortfall is not automatically intrinsic capacity. An idle rail unloading system can be starved because trains arrive late, blocked because the yard has no eligible space, held because documentation is incomplete or unavailable because it is under repair. Each mechanism can reduce observed output, but each has a different remedy. An annual effective-capacity input can absorb all of them and still produce the right historical flow. That fit would not identify which intervention will improve future service.

The diagnostic map should therefore distinguish capability, eligibility, coordination and demand. Capability asks what a ready stage can safely perform under the declared configuration. Eligibility asks whether cargo and permissions allow performance. Coordination asks whether adjacent stages and shared resources are ready together. Demand asks whether the

resulting service has a user. These categories are original analytical distinctions for the current study. They are not empirical findings or a claim that every corridor failure can be assigned cleanly to a single box.

Some mechanisms cross categories. A stockyard may block a train because its physical space is full, because segregation rules reserve that space or because a vessel cancellation changes the intended inventory pattern. The correct response depends on the cause, even though all cases present as no room. The analyst should record the immediate stop condition and the upstream reason separately. This protects the intervention from addressing the symptom while leaving the mechanism intact. It also gives managers a fairer allocation of responsibility than a single utilisation score.

A constraint can move without any asset changing. A demand shift, product mix or operating calendar can alter which resource governs the service. An intervention that improves the original stage may therefore encounter a different system by the time it is commissioned. The committee should treat the map as a living hypothesis tied to a service edition. It should request renewed evidence when the relevant conditions change rather than preserve the original bottleneck label as a permanent characteristic of the corridor.

Know when the minimum is the wrong decision model

A parallel route breaks the single-path assumption. Cargo may use several paths, each with its own capacities and shared links. The useful system constraint can then be a network cut, a dispatch allocation or an access condition rather than the smallest stage in an arbitrarily drawn route. Adding the capacities of alternatives is also unsafe when they share a bridge, border process or marine approach. A broader network method would need explicit topology and constraints. The present app should not be stretched into that role by adding imagined stages.

Finite buffers and varying calendars break the attainability assumption. Annual capacity may exist at each stage while cargo cannot be transferred at the required time because a yard fills, a gate closes or a shared handling resource is occupied. A deterministic annual minimum may then overstate realised service. It remains a useful screening upper bound, provided it is named as such. A time-based operating study should carry the chronology needed to decide whether that bound can be approached under the actual service requirement.

Quality restrictions can make gross flow a poor objective. More movement may create off-spec cargo, damage or rework that reduces accepted delivery. Legal and safety requirements can be noncompensable: a faster rate cannot offset a prohibited movement or unsafe operating configuration. A capacity calculation should therefore sit behind an eligibility gate. The model result can rank feasible options only after those conditions are reviewed, and even then the committee must decide whether increased accepted service is worth its full incremental costs.

The practical decision is to use the minimum result as a disciplined question generator. It locates which annual assumption governs the fictional flow and tests what happens when that assumption changes. It does not replace the service map, operating evidence or feasibility review. The strongest use of the bottleneck thesis is conditional: relieve the constraint that actually governs the declared service, then check the next limit and the consequences outside the annual throughput metric. An unqualified promise to optimise the bottleneck would be less defensible than this narrower, testable recommendation.

Illustrative calculation. Assumptions and sensitivity are published with the model.

Exhibit: frozen intervention sequence

All quantities below are illustrative results reproduced from the canonical engine. Other capacities, demand and operating costs remain fixed except where the named fixture explicitly changes nominal capacity.

Exhibit: frozen intervention sequence				
Study	Flow, Mt/year		Binding stage	Reading
Base chain	11.90		Rail corridor	Original annual constraint
Vessel loading only	11.90		Rail corridor	Visible expansion leaves flow unchanged
Rail only	12.00		Vessel loading	Original relief reaches the next limit
Coordinated chain	14.40		Port stockyard	Further relief shifts constraint to the yard

Constraint-validity gate

Use the annual minimum only when the intended question concerns a compatible serial annual flow. Require the model owner to identify the route, quantity conservation, demand definition and reasons annual opportunities can be combined. If an omitted queue, quality rule, shared path or calendar can reverse the decision, commission the narrower additional analysis that represents that mechanism rather than disguise it as another percentage.

1. Classify the limiting mechanism as capability, eligibility, coordination or demand.
2. Test the unchanged stages after the proposed intervention.
3. Identify which omitted condition could make the annual bound unattainable.
4. Carry the conditional result into the full service and investment decision.

A spare loader can still be commercially valuable

An operator may justify loader expansion even when annual flow is unchanged because shorter vessel service can reduce exposure to weather, free berth windows or support recovery. That is an admissible challenge to a throughput-only recommendation. It needs its own service outcome and evidence. The annual model does not calculate those effects, so neither a zero annual gain nor a positive local rate proves the commercial case. The correct response is to state the proposed benefit precisely and test it with a method capable of representing vessel windows and operating chronology.

The minimum does not identify a stochastic queue

The app has no arrival distribution, service distribution or random interruption process. A binding annual stage is not a predicted queue location, and a spare annual stage is not guaranteed to be uncongested. The model stress reduces availability assumptions deterministically; it is not a confidence interval or an event probability.

Decision implications

- Derive the bottleneck result from the declared serial assumptions.
- Test the next constraint before committing to local expansion.
- Treat non-throughput benefits as separate questions with suitable evidence.

DECISION ARCHITECTURE

04

A measurement plan can find the wrong constraint

A timestamp dictionary and evidence quality gate

Measure the transition that the decision concerns

A measurement programme should begin with the disputed decision, not with the easiest available sensor. If the question is whether a loading asset lacks intrinsic capacity, vessel turnaround alone is too broad. It includes waiting for cargo, nautical services and clearance. If the question is whether users receive a dependable end-to-end service, the loader productive rate is too narrow. The analyst should first describe the mechanism that could change the decision, then select events that distinguish it from plausible alternatives.

An event is a recorded transition, while a state is the condition between transitions. Arrival at a queue, eligibility to enter service, commencement of service, completion of service and acceptance by the next actor are distinct events. Combining them into arrival and departure may suffice to measure total residence, but it cannot diagnose why residence increased. The difference between arrival and eligibility might be documentary delay; the difference between eligibility and service start might be scheduling or resource contention. Treating every interval as equipment inefficiency creates a biased investment case.

The timestamp dictionary should state the physical location, triggering action, recording device and responsible actor for each event. It should distinguish an estimated event from an actual event and retain both rather than overwrite a forecast with an observed completion. A forecast revision is useful evidence of planning uncertainty. It is not a physical movement. If systems transmit only the latest estimate, an analyst cannot later reconstruct how much notice an operator received or whether a missed handoff arose from inaccurate information or inadequate response.

A record should also identify the cargo or service object that persists through the chain. A consignment can split between trains, a train can combine consignments and a vessel can accept several batches. The mapping between these objects is necessary to avoid duplicate quantities and missing intervals. A convenient vehicle identifier may support road-

cycle analysis while failing to reconstruct accepted cargo delivery. The unit of observation should therefore follow the decision and be reconciled with each local operational identifier.

A clean extract can still describe the wrong universe

Selection bias can enter before a timestamp is recorded. A gate log may include vehicles admitted to the terminal while excluding vehicles waiting outside. A completed-service report may exclude unfinished jobs, giving the appearance of improving average time during a growing backlog. A vessel database may record only calls that reached the berth and omit cancellations. The data can be internally clean while answering a different question from the service promise. The collection plan must state what enters the observation universe and what remains outside it.

Unfinished movements need explicit treatment. At an observation cutoff, cargo in progress has an incomplete duration rather than a short completed duration. A performance report should show the unfinished population and its elapsed waiting states. It should avoid comparing a completed-only average with a later complete cohort as though the bases were identical. This is especially important after disruption: the most delayed cases may still be unresolved while the reported completed movements recover quickly. A reassuring average can therefore coexist with persistent service failure.

Missing timestamps are not automatically zero intervals. A device failure, manual omission and unobserved event have different implications, but none justifies quietly substituting immediate movement. The analyst can mark the duration unavailable, bound it using adjacent evidence or commission a reconstruction. Each treatment should remain visible. Imputation may support a particular analysis if its rule and sensitivity are declared. It should not create an appearance that every movement was directly observed, especially where the missingness relates to difficult or congested operations.

The same discipline applies to sampling windows. A short quiet-period survey can establish a handling process but cannot represent an operating year containing maintenance, mixed cargo or concentrated demand. A campaign should deliberately include the conditions that could change the constraint classification. When those conditions cannot be observed within the available period, the report should narrow its conclusion. It may establish performance during the sampled operating envelope without asserting an annual capability or a tail-delay estimate.

Reconcile clocks, causes and quantities

Clock alignment is a substantive measurement issue at institutional boundaries. Systems can use different time zones, device clocks, manual entry conventions or batch upload times. A negative transfer duration can indicate incompatible clocks rather than physical impossibility. The dictionary should distinguish event time from record creation and transmission time. For a manually reported event, the observer and reporting delay should be known where possible. Correcting clocks should preserve the original record and the documented correction rather than erase the discrepancy.

Cause codes should be designed to support action. A generic delay code assigns no mechanism. Conversely, a long menu of mutually unclear codes invites inconsistent use. The programme can record an immediate stop condition, an attributed cause and the evidence supporting that attribution. A downstream yard full condition may be directly observed while its underlying reason remains disputed. Preserve that distinction. A model that treats an allegation about another actor as a measured cause is likely to reinforce local conflicts rather than resolve the corridor constraint.

Quantity reconciliation should track accepted movement separately from gross handling. A conveyor total can count material that later returns to the stockpile; a weighbridge total may need an agreed cargo basis; a train departure total may differ from released receipt. The end-to-end ledger should document where quantities are measured, what corrections occur and whether reconciliation differences remain unresolved. It should not force agreement by deleting the inconvenient measurement. A discrepancy is evidence of a measurement or process problem that may itself matter to the service.

The analyst should produce a lineage sample before collecting a large dataset. Select a clearly labelled fictional or consented test object and reconstruct its event sequence across the relevant systems. Check whether each handoff has an identifiable actor, timestamp basis and quantity basis. This small exercise can reveal that the intended end-to-end duration cannot be assembled from existing records. Discovering the gap early is more useful than acquiring a large extract whose columns never represented the required service transitions.

A gate extract can improve while the user waits longer

Consider a fictional appointment trial at a receiving terminal. The gate admits only vehicles with confirmed unloading readiness. Inside-terminal residence becomes shorter because ineligible vehicles remain outside or are asked to delay departure. A local report may correctly show improvement at its declared boundary. The corridor report must ask whether the driver and cargo service improved across the whole journey. Without the external and upstream holding record, it cannot distinguish removal of waiting from relocation of waiting.

The analyst should preserve the original appointment, its revisions, actual eligible arrival, gate entry and accepted receipt. A rescheduled movement is not automatically a failure, but its effect depends on the user's commitment and ability to adapt. If the carrier receives early reliable notice, upstream holding may be less burdensome than a terminal queue. If notice is late, the same policy can leave vehicles waiting on the approach. The intervention benefit therefore depends on the communication transition as well as the physical gate process.

The collection design should include movements refused entry and those not yet completed at cutoff. It should also retain the operational reason for rescheduling and the readiness confirmation that supported it. A terminal can legitimately restrict admission to protect safety or usable space. That condition belongs in the feasible operating plan rather than being described as avoidable inefficiency. The evaluation should distinguish the necessary restriction from the proposed information and appointment remedy.

A useful conclusion may be qualified: the trial improves local receipt coordination, while its end-to-end effect remains unresolved because upstream holding was not measured. That is an actionable finding. The next trial should collect the missing states and test notice quality. The team should resist both declaring universal success from the local metric and rejecting a useful local improvement merely because the wider evidence is incomplete.

Make evidence quality a decision gate

Evidence quality should affect the strength of the recommendation. A reliably measured queue at a release gate may justify an operating trial even when the wider annual capacity basis is uncertain. It cannot by itself justify a major physical expansion if the delayed population is unrepresentative or the cause unresolved. The committee should receive

the mechanism hypothesis, supporting observations, unresolved alternative explanation and proposed next action together. A dashboard without that chain of reasoning can make weak evidence look decisive.

The evidence gate can require a defined universe, consistent event semantics, acceptable reconciliation and an explanation of missingness. The threshold depends on the decision. An exploratory trial can tolerate unresolved intervals if it is designed to learn about them. A service guarantee needs a much stronger basis. The gate should state what would block the planned decision and what would merely require qualification. This prevents managers from treating every data imperfection as either immaterial or fatal.

An intervention evaluation also needs a comparison design. If a scheduling change is followed by shorter waiting while cargo mix, vessel calls and weather change at the same time, the observation does not isolate the scheduling effect. The team can still learn from the episode, but it should describe the competing explanations. A staged trial, matched operating envelope or mechanism-specific process record may provide a better basis than a broad before-and-after average. The design should be proportionate to the consequences of being wrong.

The final measurement product is a usable record of what happened, what it means and what remains unknown. It should permit an independent reader to distinguish an equipment limit from a handoff failure without trusting a local actor's preferred story. The corridor manager should be able to update the hypothesis when new events contradict it. That is the value of the timestamp dictionary: it turns a collection of local logs into a common language for service, while leaving uncertainty visible enough to influence action.

Completed event dictionary for a fictional receipt

The example defines semantics without inventing event times. Missing evidence remains explicit and directs the next collection action.

Completed event dictionary for a fictional receipt			
Event	Trigger	Recorder and basis	Missing-data action
Arrival at external queue	Vehicle reaches declared queue boundary	Gate observer; actual event time	Retain outside queue population when camera unavailable
Eligibility confirmed	Cargo and permission checks complete	Release desk; actual decision time	Mark unresolved checks rather than infer eligibility from entry
Service begins	Physical unloading starts	Equipment log; event clock	Check clock alignment with gate record
Receipt accepted	Yard accepts quantity and cargo status	Yard ledger; accepted quantity basis	Keep quality hold separate from physical unloading completion
Next handoff ready	Eligible stock assigned to loading plan	Dispatch record; actual authorisation	Retain forecast revisions separately from actual readiness

Measurement quality gate

Approve an evidence claim only after the observation universe, event dictionary and missing-data treatment match the intended service question. Require an explicit alternative explanation for an attributed bottleneck. Commission a focused collection or operating trial when a missing event can reverse the decision, instead of filling the gap with a convenient assumed duration.

1. Define the object and all inclusion rules before extracting records.
2. Reconcile clocks, quantities and completed versus unfinished movements.
3. Separate observed stop conditions from attributed causes.
4. State which unresolved defect prevents the proposed investment or service claim.

Operators need action before perfect evidence exists

Waiting for a complete dataset can itself harm service. The response is a bounded action with a learning purpose, not an unqualified investment conclusion. A reversible dispatch trial may proceed with openly incomplete evidence if it records the missing transitions and has a clear stop condition. A permanent capacity commitment needs a stronger causal basis because its cost and implementation time are harder to reverse. Evidence quality should shape the form of action and the claim made for it.

A timestamp does not explain a motive

Event records can locate delay and test a process explanation. They cannot on their own establish whether an actor acted strategically, whether a contractual obligation was breached or whether an intervention caused the observed change. Those questions need additional evidence and the appropriate review.

Evidence and scope

The WCO describes its Time Release Study as arrival-to-physical-release measurement. Its current guidance overview emphasises process mapping, sampling, segmentation, time intervals and participation across border agencies. These concepts support the collection design here; they supply no release duration for the fictional corridor.

Evidence: Time Release Study Guide

Decision implications

- Define the event before interpreting its timestamp.
- Retain unfinished and excluded movements in the measurement basis.
- Use evidence defects to select a bounded next action.

MECHANISMS AND WORKED STUDIES

05

Queues convert variability into delay

A framework for investigating eligible work, waiting and recovery across the shared service boundary

Average capacity does not describe the arrival pattern

A stage can have adequate annual capacity and still experience prolonged queues. The reason is that arrivals and service opportunities occur in time rather than as evenly spread annual quantities. A railway may deliver a concentrated batch to a yard whose handling resources are sized for a smoother pattern. A vessel can arrive while the necessary cargo is still being released. The annual model has no representation of these episodes. Its capacity margin cannot be translated directly into a promise about waiting or delivery reliability.

A queue arises when eligible work waiting for service cannot immediately access the required resource. That definition separates waiting from processing and from ineligibility. Cargo awaiting a quality release may occupy space but is not yet eligible for loading. A train awaiting an access path may be physically formed but unable to advance. A useful queue

analysis should identify the work object, the service resource, eligibility conditions and the rule determining priority. Without them, the word congestion describes a symptom rather than a mechanism.

The service pattern matters alongside its average. A resource can operate steadily, in batches or within intermittent windows. The same annual total can correspond to very different opportunities for the next stage. An inland process that releases cargo after a vessel window closes creates delay even if its annual output equals the loader's annual capacity. Conversely, a buffer can let a loader operate continuously during a window despite uneven rail arrivals. The analyst should therefore investigate matching in time rather than infer queue performance from annual equality.

This does not make annual screening useless. It tells the team where a persistent flow limitation may sit under the declared assumptions. Queue analysis then asks whether time variation and resource rules create an additional service constraint. The committee should understand the distinction between average capability and dependable completion. Increasing the former may help the latter, but the relationship depends on the actual arrival, service and interruption mechanisms. No universal queue improvement can be claimed from a capacity percentage alone.

High use leaves less room for recovery

When a stage uses much of its effective capability, it has less opportunity to absorb a concentrated arrival or recover after a short interruption. This is a qualitative consequence of finite service capacity, not a numerical waiting-time formula. The specific delay depends on chronology and the organisation of work. If arrivals can be rescheduled and the next stage has suitable stock, a heavily used resource may still support the declared service. If batches must arrive together, apparently modest average use can coexist with a difficult peak.

Utilisation needs the same basis as capacity. The canonical app compares annual common flow with effective annual capacity. That ratio describes the fictional annual balance. It does not measure the fraction of a shift spent processing, nor does it imply a stochastic traffic intensity suitable for a queueing formula. Importing the app's percentage into an unrelated waiting-time equation would join incompatible models. A quantitative queue calculation would require an explicitly implemented method, validated arrival and service assumptions, and sensitivity to their violation.

The recovery question is often more useful than a nominal utilisation target. After a disruption, how much eligible backlog can the stage clear while meeting continuing demand? Can an upstream actor restrain

arrivals without creating another service failure? Does the downstream stage have room to accept the recovered flow? A local recovery surge that overwhelms the yard simply relocates waiting. The operating plan should describe both backlog clearance and the conditions under which adjacent actors can support it.

An operator may pursue high utilisation because equipment and crews are costly to hold ready. That objective can be rational within the local contract and damaging to the corridor if it removes all recovery margin. The service design should specify what reserve capability is needed and who pays for it. Without that agreement, a local manager is rewarded for consuming the very slack that the corridor relies on during disruption. The solution is an explicit service and incentive decision, not an accusation that local managers misunderstand efficiency.

Priorities redistribute waiting

A queue discipline determines which eligible work receives service next. Arrival order, appointment, cargo urgency and commercial priority can produce different service distributions even if total annual throughput is unchanged. A priority shipment may leave quickly while ordinary cargo waits longer. A corridor report that shows only aggregate movement can miss this redistribution. The relevant service measure should therefore distinguish user classes and the commitments that justify preferential treatment.

Priority also interacts with preparation. A vessel designated urgent may not have all cargo available, leaving a loader reserved but idle while another ready vessel is excluded. A train inserted into a constrained path plan may delay several movements. The priority decision should be tested against actual readiness and system consequences. An urgent label is not a substitute for a feasible operating sequence. The coordinator should record the reason, authority and expected effect of an exceptional priority so that recurring exceptions can be reviewed.

Appointments can smooth arrivals only when the upstream actors can honour them and when service readiness is credible. A booking system that issues slots without reliable cargo eligibility, gate capacity or downstream space can move the queue outside the terminal. This may improve the terminal's own dwell measure while worsening the driver's experience and end-to-end service. The evaluation should follow the complete movement boundary, including waiting caused by rejected or revised appointments.

Fairness and capacity are separate decision dimensions. A queue rule can increase accepted service while imposing unacceptable burdens on particular users or communities. It can also preserve equal access while reducing total efficiency under certain operating conditions. The committee should state the applicable access obligations and commercial promises rather than hide the tradeoff inside a technical algorithm. The present study offers a decision framework, not a legal determination of which priority rule a particular corridor may apply.

Analyse delay with the right level of detail

The first useful delay exhibit may be a timeline rather than a simulation. It should show when eligible work arrived, when the resource became available, when processing began and when the next handoff accepted the result. Label periods of waiting, ineligibility, active work, blocking and recovery separately. This can reveal that an observed queue came from a temporary release failure rather than a persistent equipment shortage. It can also reveal that apparently isolated incidents repeatedly align with a particular scheduling boundary.

A fuller model is justified when the investment decision depends on chronology that a timeline cannot adequately compare. That model should represent the actual service objects, finite buffers, batch rules, shared resources and operating calendars that can reverse the ranking. It should preserve its inputs and be tested against suitable observations. Greater model complexity is not evidence of greater accuracy. A detailed simulation with invented arrival patterns can produce precise but unsupported service forecasts, especially when the tails of delay determine the customer's loss.

The analyst should state which evidence supports the distributional assumptions. An empirical travel-time record needs a defined route, time window and treatment of missing or unfinished journeys. A stipulated dispatch pattern should be labelled as a planning assumption. A fictional pattern should remain an illustration. Sensitivity should challenge the mechanism that matters, such as clustered arrivals or loss of a service window, rather than vary every parameter slightly and declare the conclusion robust because the same ranking survives.

Exhibit: similar waiting, different intervention

Original qualitative diagnosis. The rows are mechanisms to investigate rather than measured incident frequencies.

Exhibit: similar waiting, different intervention			
Observed state	Possible mechanism	Evidence needed	Candidate response
Ready train waits	No compatible path or priority conflict	Allocated and requested path history	Revise path and dispatch coordination
Loader idle with yard stock	Stock not eligible for the intended vessel	Grade, release and reservation ledger	Resolve release or cargo assignment
Gate queue grows	Clustered appointments exceed service window	Arrival and eligibility event sequence	Change appointments and readiness commitments
Yard cannot receive recovery surge	Downstream buffer or handling blocked	Usable space and transfer records	Coordinate backlog clearance across stages

For the fictional service, a useful diagnostic contrast is a formation waiting outside an empty yard while the receiving crew awaits a release message. The empty yard contradicts a storage-shortage explanation; it does not establish that the crew is inefficient. Reconstruct message creation, receipt and acknowledgement against cargo eligibility and crew readiness. If receipt occurs after the receiving window closes, a communication or reservation change is the candidate remedy. If receipt was timely but compatible handling was unavailable, investigate the service resource instead. If the yard was nominally empty but its available space was incompatible with the cargo, the apparent contradiction disappears. This contrast identifies the next discriminating observation without inventing a waiting-time distribution or deriving one from annual capacity.

Delay-mechanism investigation

Do not authorise an equipment expansion solely because a queue exists. First determine whether the queued work was eligible, whether service resources were ready and whether the next handoff could accept output. Use a focused timeline to choose the mechanism requiring quantitative analysis. Evaluate the remedy on the same end-to-end boundary so that displaced waiting cannot masquerade as improvement.

1. Name the queued object and eligible population.
2. Separate preparation, service, blocked time and administrative hold.
3. Reconstruct priorities and missed operating windows.
4. Test whether the proposed action changes total accepted service or merely redistributes waiting.

A queue may justify reserve capacity before detailed modelling

Reserve capability can be defensible where service failure is costly and the operating evidence shows recurrent peaks. The burden is to specify which peak or recovery duty the reserve serves, not to prove an exact universal waiting curve. A bounded operating trial can support that judgement. The decision should still include access, safety, cost and adjacent-stage readiness, because reserve equipment that cannot be deployed in the relevant window supplies no useful protection.

No waiting-time estimate is produced here

This chapter explains mechanisms and measurement choices. It does not publish a quantitative queue prediction or infer a waiting-time distribution from the annual app. A service-reliability commitment requires a suitable implemented time-based method and evidence for its assumptions.

Evidence and scope

The FHWA handbook distinguishes facility-based reliability measures from end-to-end trip measures and notes that data processing choices can change reported reliability. These scoped road-method observations support preserving the service boundary and processing rules; they establish no waiting-time distribution for this corridor.

Evidence: Application of Travel Time Data and Statistics to Travel Time Reliability Analyses

Decision implications

- A queue needs an object, eligibility rule and service resource.
- Annual utilisation cannot be substituted into an unrelated waiting-time formula.
- Evaluate whether a remedy removes waiting or moves it elsewhere.

MECHANISMS AND WORKED STUDIES

06 Buffers have location and recovery duties

A buffer location decision matrix

A buffer protects a declared interruption

A buffer should be designed around a service duty. A stockyard before the loader can allow vessel loading to continue while rail deliveries pause. A staging area before an inspection point can hold arriving vehicles without obstructing the public road. A spare locomotive can preserve movement

after a particular equipment outage. Each example stores a different resource and protects a different boundary. Calling all of them resilience obscures the conditions under which they are usable and the losses they are intended to avoid.

The buffer record should name the interruption, protected customer commitment and duty boundary. State whether protection ends after a finite campaign or must continue for later commitments, and identify replenishment where that continuing duty requires it. The question is not simply how much material fits. It is how much eligible material or resource is available for the intended duty, how quickly it can be deployed and what happens when it is exhausted. A stockyard that contains the wrong grade does not protect a vessel expecting another cargo. A spare wagon that cannot operate on the allocated route does not protect the rail service. A buffer becomes useful capacity only through a feasible deployment path.


Location changes what a buffer can protect. Stock downstream of the railway can bridge a rail interruption but cannot help clear the mine if inland storage is full. Stock upstream of the loader can protect a vessel window while exposing cargo to storage loss or quality change. A buffer at the wrong side of a closed border may preserve local output yet fail to support destination delivery. The service map should therefore show which dependency each buffer lies before or after, rather than treating total corridor storage as interchangeable.

Buffers also require handling capacity. Receiving stock, segregating it, reclaiming it and transferring it to the next stage can compete for equipment or space. A nominal storage addition can reduce usable throughput if reclaim and receipt movements conflict. The design should describe how the buffer operates during normal conditions and during recovery. A drawing that shows a larger pile without its access routes, cargo eligibility and handling sequence is a storage concept rather than a demonstrated continuity solution.

Exhibit: usable reserve and its recovery loop

Read the relationships as declared service conditions, with the limits stated in the caption.

A BUFFER SERVES A DECLARED PROTECTION DUTY



No endurance or reserve quantity is inferred from annual spare capacity.

Original conceptual diagram for a repeating protection duty. Only released, accessible stock serves the protected service. Where continuing reserve readiness is promised, the recovery loop restores stock while ongoing service is met, subject to adjacent-stage readiness. A declared finite campaign may instead end after its eligible reserve has met all promised deliveries. No buffer endurance is calculated; read the diagram alongside the completed duty record.

Usable inventory differs from total inventory

The inventory ledger should separate physically present stock from stock available for the intended service. Material may be awaiting quality release, committed to a different vessel, inaccessible beneath another pile or subject to a documentary hold. These states should remain visible instead of being counted as a single available balance. The operator needs to know which transition can release each restricted quantity and who has the authority to make that decision.

A buffer can conceal a deteriorating upstream service. The loader continues to work while the yard draws down, giving the appearance that the corridor is stable. The relevant warning is the replenishment imbalance and the remaining service duty, not merely current equipment utilisation. A useful operating report should show whether the protected function is relying on stock that the upstream chain can restore under the current conditions. Otherwise a temporary bridge becomes an unnoticed depletion process that ends in a sudden service failure.

The opposite problem is accumulation. Cargo can arrive faster than the downstream stage can accept it, filling a yard until upstream movement is blocked. Additional storage may defer the failure without changing the

persistent flow imbalance. It can still be valuable as a temporary measure if the interruption is finite and a recovery plan exists. It is less defensible as a permanent substitute for resolving a recurring downstream constraint. The design must distinguish a buffer for variability from storage used to hide an unresolved average-capacity gap.

Inventory can also carry economic and physical burdens. Cash may be tied up, material may degrade, access lanes may become crowded and grade separation may become harder to maintain. The annual app includes no inventory financing or loss mechanism. Its lower illustrative unit cost after a throughput intervention cannot establish whether additional stock is worthwhile. A buffer appraisal should connect the avoided service loss to holding, handling, quality and access consequences under the relevant operating scenario.

Recovery makes the buffer a dynamic obligation

A repeating continuity plan is incomplete when it describes depletion but not the promised recovery. After the interrupted stage restarts, it must serve ongoing flow and rebuild the protective reserve if readiness for the next interruption is part of the service duty. That requires spare usable capacity somewhere in the chain or a temporary reduction in demand served. If every stage is already committed to the normal service, the buffer cannot be restored merely by declaring the incident closed. The plan should state the recovery sequence and the user commitments that may be affected while reserves are rebuilt.

Consider an expressly fictional finite loading campaign. Its opening eligible stock covers every promised accepted delivery; no inland replenishment is available during the campaign. Loading proceeds until the final promised delivery exhausts that stock. If the declared duty ends there and includes no later vessel commitment or restored reserve obligation, the campaign can satisfy its complete service promise. Rejecting it solely for lacking replenishment would impose a repeating-service condition on a finite duty. If the promise also includes readiness for a subsequent vessel or another interruption, the same exhaustion leaves an obligation unmet. The acceptance test therefore follows the declared boundary, while preserving eligibility, handling access and all closing commitments in either case.

Recovery can shift the bottleneck. A railway returning from interruption may deliver a concentrated backlog that exceeds yard receipt capability. A loader clearing a waiting vessel may consume stock faster than the inland chain can replenish it. The corridor coordinator should manage the recovery as a shared operating state rather than instruct each actor to

maximise output independently. A surge rule must specify what the next stage can accept, how protected space is maintained and when the normal operating rhythm resumes.

The buffer itself can become unavailable during the same event it was intended to mitigate. Flooding can affect both the rail approach and adjacent storage; a power outage can stop both receiving and reclaim equipment; a documentary restriction can make physically stored cargo ineligible. This common-cause challenge matters because nominal reserve quantities may appear independent on a capacity table while sharing the decisive failure mechanism. The resilience test should ask whether the buffer survives and remains deployable under the declared interruption.

The commissioning plan should include a controlled demonstration of buffer transitions where safe and appropriate. It should test receipt, release, drawdown and the agreed closing condition, including restoration when a continuing reserve duty is promised, with clear quantity and custody records. The aim is not to manufacture a disruptive event, but to verify the interfaces that make the buffer usable. A stockpile capacity certificate cannot establish that the operator can deliver the right cargo from that stock during a restricted window. The evidence should follow the intended continuity duty.

Choose the buffer with a stopping rule

Buffer design should compare operational alternatives. Better appointment discipline, earlier documentary release or a shared dispatch plan may reduce the variability that creates the need for stock. Additional buffer capacity may still be appropriate where those measures cannot protect the customer commitment. The comparison should ask whether the proposed buffer changes the service outcome, how it is replenished and what burden it creates elsewhere. Bigger storage is not automatically more resilient if it worsens segregation or recovery logistics.

The team should state a stopping rule for adding buffer. A storage increment that no longer changes the protected service under the declared disruption should require a separate justification. Likewise, a reserve resource that cannot be deployed because of access or handling constraints should not be credited as usable redundancy. The stopping rule protects the investment from becoming a general preference for more inventory. It also helps distinguish temporary operational accommodation from a permanent structural solution.

Apply the stopping rule to a hypothetical choice between another storage bay and earlier documentary release. If existing stock cannot be loaded because release is absent, the extra bay adds gross inventory without adding eligible reserve. Earlier release may increase the usable fraction while leaving physical stock unchanged. If stock is already eligible and the protected loading campaign exhausts it before inland replenishment can resume, a storage increment may address the real duty. The missing quantity is then the campaign's eligible demand trajectory, not the annual capacity minimum. These branches tell the designer which observation separates a release remedy from a storage remedy. They leave the required stock quantity unresolved until that trajectory is supported.

A reserve approval can therefore carry a consumption disposition as well as a stock target. In the repeating-duty branch, stock consumed during protection remains a restoration obligation, and normal dispatch must leave a feasible replenishment path. In the finite-campaign branch, exhaustion can be intentional if the closing delivery and cargo-disposition obligations are satisfied. A third branch rejects reliance where the same interruption prevents stock release or loading. Recording these dispositions prevents an operator from replenishing needlessly after an intentionally terminal campaign or declaring a repeating duty complete with its protection depleted. The annual app establishes none of these inventory trajectories; the proposed disposition is an operating design to be tested separately.

Completed buffer-duty record

The illustrative record specifies the duty and exposes the missing quantities. It deliberately derives no inventory size from an annual capacity model.

Completed buffer-duty record			
Decision field	Fictional corridor record	Evidence gap	Action owner
Protected commitment	Continue eligible vessel loading during a rail interruption	Required duration and cargo sequence not frozen	Service planner defines vessel-window duty
Eligible reserve	Released cargo of required grade allocated to the vessel	Usable quantity and reclaim access unknown	Yard manager reconciles stock status
Activation	Rail replenishment fails while loading permission remains valid	Escalation trigger not demonstrated	Corridor coordinator records state transition
Recovery	Rebuild reserve without blocking yard receipt or missing next delivery	Compatible surge plan absent	Rail and yard planners agree restoration sequence
Common cause	Reserve must remain accessible during the declared interruption	Power, weather and access dependencies unresolved	Technical team challenges mitigation independence

Buffer acceptance test

Require a named interruption, protected service and duty boundary before sizing storage or spare resources. Credit only reserve that is eligible, accessible and deployable under that interruption. For a repeating protection duty, require a feasible restoration plan while continuing service is met. A finite campaign may intentionally exhaust its reserve if all promised deliveries and closing obligations are satisfied and no later reserve duty is claimed. Compare measures that reduce the disruption mechanism before treating additional stock as the only available response.

1. Locate the buffer relative to the interrupted stage and destination.
2. Identify all conditions needed to release usable reserve.
3. Plan depletion and the agreed closing condition; restore reserve when the continuing duty requires it.
4. Record the point at which more buffer stops improving the declared duty.

A larger yard may create valuable operating freedom

Additional yard space can support grade segregation, batch formation and safer handling even when it does not raise annual flow. That is a real candidate benefit, but it should be stated as such and evaluated against its particular operating requirements. The annual minimum result cannot establish or reject it. A service-specific yard plan can demonstrate why the extra space matters and which conflicts it removes, avoiding the unsupported claim that more storage universally solves the corridor bottleneck.

No reserve duration or inventory size is inferred

The canonical app has no inventory trajectory or interruption duration. Its spare annual capacity is not a stored quantity and does not imply a buffer endurance. A quantitative reserve design needs a separate time-based method, eligible inventory evidence and a sensitivity that challenges depletion and any recovery required by the declared service. The finite-duty example establishes a conditional acceptance boundary; it does not calculate a stock requirement for an actual corridor.

Decision implications

- Size buffers around a protected service and interruption.
- Count eligible deployable reserve separately from total stock.
- Restore reserve for a repeating duty; document the closing condition for a finite campaign.

MECHANISMS AND WORKED STUDIES

07 Batching changes the chain's rhythm

A batching tradeoff worksheet with transparent assumptions

A batch joins economies with commitments

Batching makes the corridor a sequence of release decisions rather than a smooth pipe. A train formation, vessel parcel or inspection lot may require a minimum eligible quantity before the next movement can begin. Larger batches can reduce repeated preparation and make handling more efficient, but they also require material to accumulate and compatible resources to be ready together. The resulting service cannot be evaluated solely by comparing annual stage rates. The timing and eligibility of the batch determine when capacity becomes useful.

The batch object should be explicit. It can be a commercial consignment, a physical loading lot, a quality sample population or an operational train formation. These objects need not coincide. A single commercial parcel may be split across trains and recombined in a stockyard; a train may contain several parcels that require separate acceptance. When the

planning system treats them as one object, a partial release can appear to complete a delivery that remains commercially unfinished. The service map should retain the links between batch identities and custody events.

A setup creates a discontinuity in productive operation. Equipment may need repositioning, cleaning, inspection or a configuration change before the next batch. A larger batch can spread that preparation across more cargo while reducing flexibility. If the next vessel requires a different grade, a long current run may delay its preparation. The apparent productivity gain therefore needs to be evaluated against the relevant user commitments. There is no universally best batch size independent of product mix, service windows and the costs of waiting or changeover.

A batch also reserves resources. A train formation can occupy a siding, wagons and a loading crew while awaiting final release. A vessel parcel can reserve usable stock that another vessel might otherwise receive. These reservations can prevent local utilisation even when physical capacity exists. The diagnostic record should distinguish reserved, ready and actively used resources. Otherwise managers may respond to low productive time by adding assets when the actual constraint is how commitments are assembled and released.

Synchronise batches across unequal rhythms

The chain's rhythms differ. Producer release may follow extraction and quality control; railway service follows path availability and fleet circulation; vessel service follows calls and marine access. An annual capacity model assumes these rhythms can be reconciled sufficiently to achieve its common flow. In a real operation, mismatches can create repeated waiting. The planning team should identify which stage sets the cadence for each service class and how the other stages prepare eligible work to meet it.

A railway batch that arrives after a vessel cargo cutoff may be timely under the rail plan and late under the delivery promise. Conversely, accelerating that train may occupy a path or yard needed for another committed service. The corridor coordinator should compare the end-to-end consequences of changes to release times. A local improvement is insufficient if it merely transfers the missed commitment to another user. Batch planning requires a common calendar of readiness and acceptance rather than separate local departure targets.

Consolidation can improve equipment use while increasing the time an early consignment waits for other cargo. The balance should be stated in service terms. Does the user accept a longer accumulation interval in

exchange for a lower operating cost or more dependable departure? If different users have different commitments, the planner may need several service classes rather than force all material into the same batch policy. An aggregate throughput calculation cannot resolve that commercial choice because it contains no distribution of consignment completion times.

Segregation can constrain consolidation. Cargo of incompatible grades or release statuses may require separate piles, loading sequences or transport arrangements. A nominally full train or vessel parcel may be unusable if the combination violates an acceptance requirement. The team should record which compatibility rules are mandatory and which are operating preferences. Mandatory conditions should form the feasible set before batch efficiency is compared. The convenience of a larger batch cannot compensate for an unacceptable cargo mix.

A dispatch rule needs readiness evidence

Dispatch should rely on readiness at the relevant downstream boundary. A physically formed train is not ready for an end-to-end service if the unloading facility cannot receive its cargo or the yard cannot preserve eligibility. A loader plan is not ready if it assumes an unconfirmed vessel window. The dispatch record should name the conditions that must hold and the actor confirming each condition. This prevents a local departure from being counted as progress when it merely creates a remote queue.

Forecast readiness and confirmed readiness should remain separate. Planners need forecasts to allocate paths and resources before every condition is final. Operators need explicit confirmation before movement that cannot safely or economically be recalled. A useful plan can tolerate uncertainty by specifying the last reversible decision and the evidence required beyond it. The corridor should not resolve the tension by labelling every forecast green or by refusing to plan until all future events are certain.

Batch priority should have an expiry or review trigger. A cargo initially selected for an urgent vessel can lose that role when the vessel window changes. Keeping its reservation indefinitely may block other ready work. The coordinator should record who can release the reservation, whether commercial consent is required and how the revised allocation affects commitments. The operational record then preserves the reason for the decision rather than leaving later reviewers to infer intent from an unexplained sequence.

Exceptional dispatch also needs a system consequence check. Sending a partial train can relieve a mine stock constraint while reducing fleet productivity or consuming a scarce path. Holding it can preserve transport efficiency while worsening a producer commitment. The correct choice depends on the service objective and agreed authority. The annual app cannot adjudicate the tradeoff. It can show whether the chain has an annual capacity gap, while the batch record supplies the local information needed to decide how to operate within that gap.

Evaluate a batch policy as an operating intervention

A batch-policy trial should state the changed rule rather than merely announce improved coordination. It might alter release eligibility, reservation expiry, appointment sequence or the point at which downstream readiness is confirmed. Hold the relevant operating envelope as comparable as possible and record competing changes in cargo mix or demand. Evaluate accepted service and resource burden together. A policy that accelerates the preferred parcel while systematically delaying other users should not be reported as a general corridor improvement.

The trial should also test recovery. A normal batch plan can become unsuitable after interruption when backlog and current demand compete for the same slots. Larger recovery batches may clear material efficiently while blocking urgent smaller commitments. A recovery rule should explain who can depart from the normal policy and when the normal rhythm resumes. This is a governed decision rather than a parameter that the annual capacity calculation can choose automatically.

The committee should distinguish reversible operating choices from physical commitments. A release rule can often be tested and revised before equipment is procured. A larger siding, stockyard or loading system may create more batch flexibility but at a cost and with access requirements. The operating trial can reveal whether those physical additions address a persistent incompatibility or merely accommodate a poorly defined dispatch policy. It also generates evidence for the design brief by showing which batch states actually consume scarce resources.

A completed qualitative policy contrast can distinguish early commitment from late confirmation. In a hypothetical early-reservation rule, the formation retains its place while downstream readiness is forecast; a later failure can leave the slot unusable. In a confirmation rule, the slot is released when readiness is unconfirmed and reallocated to another eligible formation. The latter can reduce stranded reservations while

disadvantaging cargo that needs longer preparation. The trial should therefore retain expired reservations, reallocated slots, displaced consignments and accepted receipts, rather than just the successful departures. The useful inference is conditional: prefer the confirmation rule only if reduced stranded opportunity exceeds its preparation and access consequences under the agreed service priorities. No optimal batch size follows from the annual cases.

Completed qualitative batch-policy comparison

The fictional comparison records why a policy could help and where evidence remains absent. It calculates no optimal batch size.

Completed qualitative batch-policy comparison			
Policy	Intended effect	Possible burden	Evidence before adoption
Consolidate ready cargo into larger rail formations	Reduce repeated formation and loading preparation	Early cargo waits and siding reservation grows	Eligible batch arrivals, setup events and missed service commitments
Release partial formation for an urgent vessel	Meet a specific destination window	Consumes path and fleet cycle with less cargo	Window confirmation, path permission and alternative-user impact
Reserve yard stock only after vessel readiness confirmation	Reduce unusable cargo reservation	Less planning certainty for the loader	Forecast revisions, confirmation timing and release authority
Expire priority after a changed vessel window	Restore access for other ready cargo	Commercial consent or reassignment dispute	Priority basis, reservation record and accountable approval

Batch-policy trial gate

Select a policy only after defining the batch object, mandatory compatibility rules and downstream readiness evidence. Identify the last reversible dispatch point. Trial the rule against accepted service for all affected users, not only local equipment productivity. Revisit it during recovery or when the product mix changes.

1. Map commercial parcels to physical lots and formations.
2. Record setup, reservation and eligibility transitions.
3. State who confirms downstream readiness and who can release a reservation.
4. Compare service gains with waiting transferred to other users.

Large batches can be essential to a competitive service

A bulk service may depend on train and vessel scale, making smaller batches commercially unattractive. The paper does not recommend smaller batches universally. It recommends exposing the waiting, reservation and compatibility consequences of the chosen scale. A large batch can be the right decision when the service commitments and supporting storage are designed around it. Its scale still does not prove that the adjacent stages can prepare and accept it in the relevant windows.

No batch optimum is claimed

There is no frozen setup-cost or dispatch model in the current app. The comparison is an original operating framework and does not publish a calculated optimum. Numerical batch sizing should use a suitable implemented method and declared inputs rather than infer time patterns from annual capacities.

Decision implications

- Define which batch is physical, commercial and eligible.
- Coordinate readiness at the next boundary before irreversible dispatch.
- Evaluate batch efficiency against every affected service commitment.

MECHANISMS AND WORKED STUDIES

08

A port expansion can move the bottleneck inland

A canonical corridor app intervention comparison

The visible asset is not necessarily the annual constraint

Port investment often attracts attention because its assets are concentrated, observable and easy to illustrate. A loader, berth or crane can be photographed and its local rate can be tested. Inland interfaces may be dispersed across paths, loading points, border processes and maintenance arrangements. Visibility affects how a problem is described, but it does not determine what constrains the declared service. The system decision should compare effective annual capacities and the relevant time-based operating conditions before selecting the most tangible expansion.

The fictional base case makes this distinction deliberately clear. Rail is the annual binding stage while the loading asset has a slightly greater effective capacity. Increasing loading capability alone leaves the common flow unchanged. This is not a criticism of ports or a claim that inland assets are

generally the constraint. It is a controlled counterexample to an asset-first inference. The paper's title describes a method of reasoning, not a geographical generalisation about where bottlenecks occur.

A port itself is a chain of functions. Marine access, berth occupation, cargo handling, storage, release and landside transfer can each limit a particular service. An expanded berth can be unusable when its approach, cargo supply or adjacent handling is constrained. More storage can be valuable for segregation without increasing the loader rate. A faster loader can reduce service time for a ready vessel without increasing annual corridor flow. These outcomes should be named separately so that a legitimate local benefit is neither exaggerated nor dismissed.

The investment question should therefore specify the margin being improved. Is the proposal intended to increase accepted annual quantity, reduce vessel exposure, improve recovery, serve a new cargo class or preserve an access obligation? The canonical model answers only the first question under its restricted annual assumptions. The other objectives require evidence and methods representing their mechanisms. A committee can approve a port intervention for a non-throughput reason, but it should not justify that decision with a throughput number that the model does not produce.

Compare interventions on one declared basis

Illustrative calculation. Assumptions and sensitivity are published with the model.

The base study calculates 11.90 Mt/year. The vessel-loading intervention increases the nominal loading input to 20.00 Mt/year while preserving its illustrative availability and the rest of the chain. Calculated flow remains 11.90 Mt/year. The unchanged Rail corridor still sets the annual limit. The case contains no capital cost or delivery schedule, so the result should be read as physical sensitivity, not a project recommendation.

Illustrative. The rail-only study increases nominal rail capacity to 20.00 Mt/year. The calculated flow becomes 12.00 Mt/year and Vessel loading becomes binding. The small change in delivered flow relative to the nominal intervention illustrates a near-adjacent constraint. An analyst who presented the added nominal rail rate as added corridor flow would overstate the result. The relevant quantity is what survives the minimum across all unchanged stages.

Illustrative. The coordinated study changes rail and loading nominal capacities together. It calculates 14.40 Mt/year and identifies Port stockyard as the next constraint. This case supplies a different intervention logic: joint relief can achieve a result that neither isolated visible expansion

nor the original-stage improvement reaches alone. It still does not establish that simultaneous investment is preferable to a staged sequence. Timing, feasibility, incremental cost and learning value remain unmodelled.

The comparison preserves one cargo quantity and one demand assumption. In an actual proposal, an expansion may introduce a different commodity, customer or handling method. That can be a valid business change, but it is no longer the same intervention comparison. The analyst should write a revised service and cost basis, not attribute every change in output to added capacity. Controlled comparisons are valuable because they separate the mechanism under test from unrelated improvements that would otherwise make the conclusion impossible to challenge.

Identify the cost of moving the constraint

Moving a constraint can create new burdens. More rail delivery can require additional yard receipt, segregation and reclaim readiness. A greater loader rate can change stock drawdown and the relationship between rail batches and vessel windows. A port expansion that raises the annual bound may therefore need coordinated operating changes before that bound becomes usable. The committee should ask which adjacent interface becomes more demanding and who is responsible for preparing it.

Illustrative calculation. Assumptions and sensitivity are published with the model.

The annual cost comparison reveals fixed-cost dilution, but it excludes the incremental obligations created by intervention. The base operating unit cost is 23.73 USD/t, while the coordinated case is 22.39 USD/t under unchanged cost inputs. This decline is mechanically produced by higher common flow. It is not a forecast tariff or an estimate of investment return. Any actual expansion appraisal must add the new capital and operating consequences rather than treat the illustrative denominator improvement as the entire economic benefit.

The constraint may also move outside the physical chain. Greater output can exceed contracted demand or require permissions that have not been secured. A port capable of loading more cargo is not useful if the customer cannot receive it or if a marine service is unavailable in the relevant window. The expanded service should be tested against eligible demand, access, quality and safety. The annual model's demand input is a bound chosen for the fictional comparison rather than evidence that incremental cargo can be sold or accepted.

A staged plan can use the next constraint as a learning point. Improve the original stage through a bounded operating measure, observe the response at the yard and loader, then revise the physical design if the next limit is confirmed. This sequence may reduce irreversible commitment, but it is not universally preferable. A construction window, shared mobilisation or service obligation can make coordinated works more efficient. The committee should compare staged and joint alternatives on explicit costs and risks, using the annual result as a constraint map rather than a complete investment selector.

A port decision needs an interface acceptance plan

The design brief should translate the intended benefit into interface conditions. If added loading capability is justified by shorter vessel service, define cargo readiness, access windows and the measurement boundary for service time. If it is justified by higher annual delivered quantity, define the inland replenishment and yard-release basis that must support it. If it is justified by resilience, define the interruption and recovery duty. These statements allow procurement and commissioning evidence to address the actual service rather than a generic equipment rate.

An acceptance test at the machine can verify local performance under the agreed conditions. A corridor acceptance test must also verify eligible cargo supply, coordinated receipt and loading, documentary release and destination transfer. Both are useful, but they establish different properties. The sponsor should preserve the local equipment warranty while adding a separate system test. Asking a supplier to guarantee uncontrolled inland service is likely to create an unworkable contract; omitting the system test leaves the sponsor without evidence that the investment serves its purpose.

The predicted stockyard constraint provides a concrete receiving test. Ask the fictional yard team to trace an eligible rail arrival through unloading, segregation, stock location and the next loading instruction under the proposed coordinated duty. A failure before unloading points to receiving access or handling; a failure after unloading points to usable storage, cargo compatibility or release to the loader. If both transitions succeed but formations cannot leave, the apparently downstream problem also consumes rail circulation. This decomposition turns a stockyard label into candidate duties that can be challenged. It avoids authorising a larger yard before determining whether the scarce function is storage volume, handling opportunity or resource release.

Separate the port recommendation into a flow proposition and a window proposition. The loading-only fictional case fails the annual-flow proposition because inland rail remains binding. It can still motivate a window study if faster loading would release a berth during a scarce access period. The coordinated case passes the model's annual-flow comparison but may fail that window study if eligible stock does not reach the loader in time. These conclusions can coexist without altering the frozen results. A port brief should identify which proposition justifies its design, the interface test supporting it and the incremental obligations omitted by the annual operating-cost allocation. A window benefit should never be counted again as extra annual tonnes without its own supporting mechanism.

Illustrative calculation. Assumptions and sensitivity are published with the model.

Exhibit: physical gain and economic omission

Illustrative canonical operating outputs. The intervention investment and any incremental operating cost are absent, so the table cannot rank net investment value.

Exhibit: physical gain and economic omission			
Case	Flow, Mt/year	Operating cost, USD/t	Omitted decision input
Base	11.90	23.73	Actual demand and operating evidence
Loading only	11.90	23.73	Added loading capital and window benefit
Rail only	12.00	23.67	Rail feasibility and next-stage readiness
Coordinated	14.40	22.39	Joint investment, timing and operating obligations

Port intervention brief

Require the sponsor to name the benefit as annual quantity, vessel-window performance, resilience or another defined service. Select evidence and an acceptance boundary for that benefit. Where the annual minimum is relevant, compare the whole chain after the port change and identify the next constraint. Return for revision an economic recommendation that uses unchanged illustrative costs as though they included the expansion investment.

1. State the user outcome and boundary.
2. Hold the comparison basis constant or document its change.
3. Review adjacent-stage readiness and permissions.
4. Separate local equipment performance from whole-chain acceptance.

Marine service benefits can justify a throughput-neutral expansion

A loading improvement may reduce a ready vessel's berth occupation or support operation in shorter safe windows. Those benefits can matter even when the annual chain result is unchanged. The annual comparison therefore cannot establish that the intervention has no value. It establishes only that its justification cannot be an unqualified increase in the modelled common annual flow. A suitable time-based and economic study should evaluate the marine service proposition separately.

Operating cost dilution is not a capital return

The canonical engine contains annual fixed and unit variable operating inputs, with no capital schedule, financing, discounting or residual value. No investment return, bankability or transaction recommendation follows from the displayed operating unit costs. Physical feasibility and jurisdiction-specific obligations also remain outside the computation.

Evidence and scope

The IMO Just in Time portal links arrival planning to coordinated readiness of berth, fairway and nautical services. It distinguishes relevant operational timestamps and data exchange from physical asset capability. This supports the interface questions here, without importing any reported emissions saving or port-specific performance estimate.

Evidence: Just in Time Portal

Decision implications

- Judge the port proposal at its intended service boundary.
- Show the next constraint after every capacity intervention.
- Add incremental investment evidence before selecting an economic alternative.

MECHANISMS AND WORKED STUDIES

09

Rail paths differ from rolling-stock capacity

A rail operating architecture and constraint test

A path is a scheduled right to move

Rail capacity cannot be described adequately by the number of locomotives or wagons. A train requires a compatible path through the infrastructure, a feasible formation, terminal access and readiness at both ends. A path has timing and operating characteristics; it is not simply an annual tonne allowance. Several trains can compete for the same infrastructure resource, and passenger or maintenance requirements can

shape the usable pattern. The annual effective-capacity input should be derived from that operating architecture when applied to an actual corridor.

Infrastructure capability and access permission are distinct. A line may physically accommodate a train while the operator lacks the appropriate right or the desired timetable slot. Conversely, an allocated path may be unusable because terminal turnaround, traction or cargo readiness fails. The service analysis should preserve both conditions. Adding rolling stock cannot create a path automatically, and obtaining a path cannot make an incompatible or unready consist deliver cargo. A fleet procurement case should therefore identify the movement opportunity it will actually use.

The system boundary includes terminal approaches and release. A loaded train can occupy infrastructure while awaiting yard receipt, reducing the opportunity for other movements. A locomotive can remain tied to a train whose unloading is delayed. These are circulation consequences, not merely terminal statistics. The diagnostic record should track the train's physical movement and the resource cycle that returns locomotives and wagons to service. An annual capacity estimate that assumes prompt turnaround needs evidence at the destination, not only a route speed assumption.

Rail-specific safety and access rules depend on the jurisdiction and infrastructure manager. This paper uses publicly available institutional guidance as scoped examples of the architecture, without asserting that those rules govern an unnamed corridor. A real application should obtain the relevant operating requirements and agreements. The analyst's task is to show where they enter the feasible service and how they constrain the capacity basis. It is not to replace an access review with a generic capacity multiplier.

The fleet is a circulating stock

Wagons deliver service through repeated cycles. Their cycle includes preparation, loading, movement, waiting, unloading, return and maintenance. A nominal fleet count is therefore a stock from which a flow is derived under a calendar and turnaround basis. A slower destination release can reduce the service supplied by the same fleet. A shorter empty return can improve fleet productivity while leaving the infrastructure path constraint unchanged. The intervention diagnosis should identify which part of the cycle limits the intended output.

The same logic applies to locomotives and specialised equipment, but their cycles may not align with wagon cycles. A locomotive can be detached, substituted or assigned to another service while wagons remain at a terminal. Treating an entire formation as a single permanently coupled resource may overstate or understate the actual requirement. The operating architecture should state which resources circulate together, where substitution is permitted and what evidence demonstrates the relevant turnaround.

Payload depends on the permissible and practical configuration. Train length, axle load, locomotive performance, braking, route geometry and terminal handling can affect what an operating formation can carry. The paper does not quantify these relationships because the app does not implement them. For an actual project, the capacity basis should retain the technical assumptions and the actor responsible for validating them. A headline annual rate derived from an unspecified payload and cycle time is not a defensible basis for procurement.

Maintenance reserve should also have a defined duty. A fleet availability assumption may already include planned and unplanned unavailability, while a separate spare count may protect a particular interruption. Combining both without reconciliation can double-count protection. Conversely, a nominal spare can be useless if it lacks compatible equipment, certification or a deployment arrangement. The record should distinguish mechanically ready, authorised and operationally deployable resources, using the same service boundary that the fleet expansion is intended to improve.

Locate the interface that consumes the cycle

The destination terminal often determines how quickly rolling stock returns to the corridor. Waiting for a free receipt track, cargo release or an unloading resource can tie up wagons even when the line itself has spare capability. A terminal investment can therefore improve fleet circulation without adding track capacity. That proposition needs a cycle record showing the delayed transition and the effect of the proposed remedy. It should not be inferred from a low fleet-utilisation statistic alone.

A rail path improvement can likewise fail to increase accepted service if the producer cannot form eligible cargo in time. The operator may then run underfilled trains, miss allocated slots or hold equipment waiting for release. The appropriate intervention could be a cargo preparation rule rather than more trains. The diagnosis should ask whether a path was

requested, allocated, available and actually used, and whether unused opportunity came from cargo readiness, equipment or another operating condition.

Direction and timing matter. Loaded and empty trains may use different resources or compete differently with other traffic. A concentrated loading programme can create a return imbalance that affects the next dispatch period. A terminal closure can block receipt while empty equipment is still needed inland. The annual serial model removes these time and direction details. It can screen an annual rail-capacity hypothesis, but a timetable and circulation analysis is needed when those details determine whether an expansion is usable.

The planning team should identify shared dependencies. Alternative rail operators may rely on the same terminal, bridge, dispatch function or maintenance window. Their nominal capacities cannot be added as independent resilience. Likewise, spare locomotives cannot compensate for an unavailable infrastructure segment. A useful operating architecture shows both the movement chain and resource circulation, with shared constraints marked explicitly. This prevents a fleet solution from being credited against a failure mechanism it cannot address.

Test a rail intervention at the corridor boundary

Illustrative calculation. Assumptions and sensitivity are published with the model.

The fictional rail-only case changes the annual nominal rail input and produces 12.00 Mt/year. The next binding stage is Vessel loading. This controlled result tells the reader that relieving the original annual rail limit reaches another constraint. It does not show that additional train paths, fleet cycles or terminal arrangements are technically feasible. Those would be necessary inputs to justify a real revision of the annual capacity basis.

A rail intervention brief should therefore name the mechanism: additional paths, greater usable payload, improved circulation, higher readiness or better destination release. Each mechanism needs different evidence. A path case needs a compatible operating plan and access basis. A fleet case needs a cycle record. A terminal case needs handoff and resource evidence. Describing all of them as more rail capacity conceals the difference between a physical asset, an operating change and a permission.

The trial should follow delivered cargo and fleet return rather than count departures alone. A successful dispatch that leaves wagons immobilised at destination can weaken the next service cycle. The evaluation should record whether the intended next-stage acceptance occurred and whether the resource became available again under the declared calendar. This is

especially important during ramp-up, when a first movement may be easy to stage while a sustained repeating service exposes unresolved circulation dependencies.

A hypothetical fleet proposal can be rejected without rejecting a path proposal. Suppose formations are ready but scheduled paths expire while the destination cannot release the previous formation. Additional wagons increase the waiting stock; additional paths may remain unused. The first candidate action is then destination release or its coordination with the return movement. If destination release is prompt but eligible formations repeatedly miss available departures because locomotives are committed elsewhere, traction assignment becomes the relevant branch. Preserve unused paths and immobilised formations in the same cycle account. Counting only successful departures would conceal both failures and could make either procurement look effective. The annual rail input should change only after the branch that creates usable recurring movement has been demonstrated.

Exhibit: rail capability has separate dependencies

Original operating architecture. No actual network or legal access conclusion is represented.

Exhibit: rail capability has separate dependencies			
Resource or condition	What it permits	Failure that it cannot repair	Evidence owner
Allocated compatible path	Timed infrastructure movement	Unready cargo or unavailable train formation	Infrastructure manager and planner
Deployable locomotive	Traction for permitted formation	Closed route or unavailable destination	Rail operator
Compatible wagon fleet	Cargo carriage through repeat cycles	Blocked unloading or ineligible cargo	Fleet and terminal managers
Destination release	Unloading and resource return	Missing inbound access or unsafe formation	Receiving terminal
Operating rules and maintenance access	Safe feasible timetable basis	Commercial demand shortfall	Relevant jurisdiction and infrastructure actors

Rail capacity basis test

Require the proposed annual rate to trace to a usable path plan, permitted formation, resource cycle and destination acceptance basis. Identify the mechanism changed by the intervention. Confirm that the unchanged stages can accept the resulting service and that recurring operation, not merely a first train, is represented.

1. Separate access rights from infrastructure capability.
2. Trace cargo movement and resource return as different records.
3. Identify the terminal transition that consumes the cycle.
4. Validate the resulting annual basis before changing the corridor input.

Fleet procurement can be justified before every path is final

A long procurement lead time may require a conditional fleet decision while the timetable develops. That can be a rational risk-bearing choice. The committee should record the unresolved access and terminal conditions, the ability to alter or defer procurement and the consequences if the assumed paths are not secured. The annual model cannot remove those risks; it can only show what the declared rail rate would mean for the rest of the chain.

No train-path or fleet-cycle calculation is implemented

The Corridor lab accepts an annual rail stage like any other serial stage. It does not calculate timetable conflicts, rolling-stock requirements, permitted payload or technical route compatibility. This chapter provides an architecture and evidence test, not a reproduced rail engineering design.

Evidence and scope

Network Rail's statement separates infrastructure characteristics, access conditions, capacity allocation, operations and service facilities. It points to engineering access and timetable planning rules as distinct controlled documents. This illustrates the layered rail capacity basis in Great Britain; no local rule or path capacity is inferred for the fictional corridor.

Evidence: Network Statement for the timetable year

Decision implications

- Treat paths, formations, fleet circulation and terminal acceptance as distinct conditions.
- Trace a rail intervention to the mechanism it actually changes.
- Verify sustained cycles rather than only first departures.

MECHANISMS AND WORKED STUDIES

10

Road capacity is not a truck count

A truck-cycle model and evidence collection record

A truck supplies a repeated service cycle

Road haulage capacity is a flow produced by vehicles, drivers, roads and terminals operating together. A fleet count is an input stock, not a delivery rate. The useful cycle begins when the vehicle is eligible to receive cargo and ends when it can perform the next compatible service. It includes loading preparation, movement, queues, inspection, unloading, return and any required rest or maintenance. An analysis that counts only loaded travel omits the states that often determine whether another trip is possible.

Vehicle payload should match the cargo and route basis. Gross vehicle capacity is not necessarily the accepted cargo quantity, and a nominal payload may be restricted by route, vehicle configuration, cargo characteristics or applicable rules. The team should identify the relevant technical and legal constraints without substituting a generic industry figure. This paper publishes no truck payload or cycle duration because the canonical app does not calculate them. A real road-stage input requires a validated derivation rather than a plausible-looking fleet multiplication.

Drivers and vehicles may have different availability patterns. A ready vehicle does not imply an available qualified driver, and a driver roster does not repair an unavailable vehicle. Applicable working-time, rest and safety requirements belong in the feasible operating envelope. The analysis should obtain the relevant jurisdiction-specific basis rather than assume that every hour on a calendar can become productive driving. Increasing dispatch pressure without that review can convert an apparent capacity solution into an unacceptable operating practice.

Terminal readiness also enters the cycle. A loading queue can immobilise vehicles before travel begins; a receiving queue can delay return; a documentary hold can occupy cargo and equipment without productive movement. A road expansion proposal should state which transition it improves. More trucks may increase the external queue when the terminal remains constrained. A faster loading process may improve fleet

circulation without altering the road. Treating all shortfall as insufficient fleet size prevents the committee from comparing those alternatives intelligently.

Road condition changes the operating envelope

Travel capability depends on the route's physical and operating conditions. Road surface, gradients, crossings, traffic interaction and restrictions can affect a permitted vehicle configuration and the consistency of travel. Seasonal conditions can alter access or require a different operating plan. The capacity basis should name the route and the relevant envelope, including the conditions under which movement must stop. It should not treat an annual road-stage capacity as a promise that every journey is feasible in every condition.

A corridor may use public infrastructure shared with communities and other traffic. Its operating plan should account for that shared use and the applicable obligations, rather than define success only as moving more industrial cargo. A queue displaced to a public approach can create a burden outside the terminal's own measurement boundary. A dispatch improvement should therefore be tested against the route and user interfaces that it affects. Local gate performance alone may conceal the external consequences.

An alternative route needs its own compatibility and permission evidence. A line on a map does not establish that the loaded vehicle can use it, that the receiving terminal is reachable or that the movement is authorised. Travel time on an unloaded reconnaissance trip does not demonstrate loaded service. The resilience record should identify which parts of the alternative have been tested and which remain uncertain. Otherwise a nominal detour can be credited as reserve capacity without a feasible deployment path.

Road conditions can also be endogenous to the service. Greater industrial use can increase maintenance needs, change local traffic patterns or affect access arrangements. The expansion basis should identify who bears those consequences and who can authorise the resulting operating regime. The annual app's cost inputs do not calculate road deterioration, public impacts or user-specific delay. Those matters require additional evidence and review, but they should be visible in the service decision rather than silently assumed away.

Observe the cycle without mistaking speed for service

A vehicle trace can show movement while missing eligibility and custody. A truck can remain stationary for a permitted rest, a mechanical failure, a queue or a cargo hold. The same low speed has different meanings. The cycle record should connect location and movement to loading, release, receiving and resource-status events. An analyst should not infer a cause from a trace alone when the operational record can distinguish those states.

Travel-time reliability should use a declared trip boundary. A report of road segment times does not establish the completed industrial delivery cycle. The full cycle includes terminal and administrative intervals, while a specific segment measure remains useful for diagnosing road conditions. Keep both measures if they serve different decisions. Combining them without explaining their objects and denominators can create an apparent contradiction between a road that performs well and a delivery service that remains unreliable.

A vehicle sample can also be unrepresentative. The most complete traces may belong to newer equipment or preferred contractors, while difficult movements have missing records. Completed trips may exclude vehicles still delayed at the cutoff. The evidence plan should preserve those exclusions and classify them before reporting a representative result. A neat route average based on convenient vehicles should not become a corridor-wide capacity estimate merely because the dataset is large.

An operating trial should record dispatch and arrival management as well as travel. If appointments change, the evaluation should include vehicles that wait outside the gate or are rescheduled before departure. If terminal service improves, track whether vehicles actually return sooner and can undertake the next eligible movement. A shorter local dwell period may be valuable, but it is not equivalent to a greater sustained fleet flow unless the rest of the cycle permits that response.

Select the intervention that changes the limiting state

A road intervention diagnosis should ask whether the service is constrained by vehicle circulation, terminal handling, route access, driver readiness or eligible demand. Each answer changes the candidate action. More vehicles address a circulation-resource limit only when drivers, terminals and routes can support them. Appointment changes address

arrival mismatch only when service readiness is credible. Road works address a route condition only when that condition governs the declared service and the improved route remains accessible.

The intervention should carry a falsifier. If added vehicles increase queue length without increasing accepted receipt or resource return, fleet quantity has not demonstrated the expected system gain. If appointment discipline reduces terminal waiting but creates longer upstream holding, the end-to-end effect needs reassessment. If a road improvement increases permissible movement but demand remains lower, the physical benefit may not create additional service. These are testable mechanisms rather than generic statements that every improvement will optimise logistics.

A road-stage annual capacity can enter the canonical app after the separate cycle and operating analysis produces a coherent annual bound. The app then compares that bound with other serial stages. It cannot verify the derivation or simulate traffic, driver rosters and terminal queues. The practitioner should therefore retain the underlying record and review it when the route, vehicle mix or service promise changes. A copied annual input without its basis loses the information needed to judge whether the next comparison remains valid.

Completed truck-cycle evidence record

This fictional collection plan contains no invented cycle observations. It identifies the evidence needed before deriving an annual road-stage rate.

Completed truck-cycle evidence record			
Cycle state	Boundary definition	Current evidence	Next action
Loading readiness	Vehicle, cargo and permission jointly ready	Not measured in the fictional study	Link dispatch, release and vehicle-status records
Loaded travel	Departure from load point to declared receiving approach	No route trace or loaded envelope supplied	Observe compatible loaded route under declared conditions
Receiving wait	Eligible arrival to start of unloading	Outside and inside queues unresolved	Record both queues and the receiving resource state
Return readiness	Vehicle and driver eligible for next assignment	Maintenance and rest basis absent	Obtain relevant operating rules and circulation events
Annual bound	Repeatable cycles under the stated calendar	Not derivable from the annual app alone	Validate a separate cycle method before entering capacity

A useful truck diagnosis follows the resource that is actually scarce. In a hypothetical shift, vehicles return promptly but drivers cannot take their next admissible assignment; increasing the fleet does not relieve that driver-state limit. In another shift, drivers are ready but vehicles remain occupied at receipt; the candidate remedy lies at the terminal. A route restriction can defeat both branches even when vehicles and drivers are available. The collection instrument should retain these incomplete cycles and pair each with the next permissible assignment, rather than infer readiness from a vehicle arriving home. Comparing the branches identifies whether the design needs more circulation resources, better receipt or a different operating calendar, while preserving the safety and access conditions already established.

Fleet expansion challenge

Require the fleet proposal to identify the cycle state that limits accepted service. Confirm driver, route and terminal readiness before crediting additional vehicles. Evaluate any operating trial on complete cycles and end-to-end delivery, including displaced queues. Preserve noncompensable safety and access conditions before comparing physical output.

1. Define cargo, vehicle and trip boundaries.
2. Reconstruct waiting, movement, service and return eligibility.
3. Test the limiting state against alternative explanations.
4. Trace the resulting annual capacity assumption into the corridor model.

A fleet reserve can protect continuity without increasing average flow

A reserve vehicle may be justified by replacement duty during breakdown or maintenance rather than by higher annual throughput. That is an admissible service objective. The reserve needs compatible equipment, available drivers and a deployment process under the interruption it addresses. The annual serial calculation cannot value that continuity duty and should not be used to reject it merely because the base throughput remains unchanged.

The road architecture is conceptual

No quantitative truck-cycle, traffic or driver-scheduling method is implemented by the Corridor lab. This chapter supplies a measurement and decision architecture. It does not establish route capacity, lawful operating hours or a required fleet count for an actual project.

Decision implications

- Derive a fleet flow from complete eligible cycles.
- Include terminal, driver and access states alongside travel.
- Test the mechanism before buying more circulating resources.

MECHANISMS AND WORKED STUDIES

11

Maintenance introduces calendar dependence

A maintenance-window and recovery planning framework

Downtime is placed on a calendar

Maintenance affects capacity through timing, duration and the state in which equipment returns. An annual availability factor can represent a declared average reduction, but it cannot show whether the lost opportunities coincide with customer windows or whether adjacent stages can use the remaining opportunities. The same annual factor can describe a planned closure in a quiet period or repeated failures during critical service. These patterns can have very different consequences for accepted delivery.

Planned maintenance is a deliberate operating intervention. Its purpose is to preserve capability and control failure exposure, not simply to subtract productive hours. A capacity appraisal should therefore avoid treating maintenance reduction as an automatic efficiency gain. Deferring work can raise apparent short-term availability while weakening future reliability or violating the approved operating basis. The relevant decision concerns the service and lifecycle consequences of the maintenance plan, with technical and safety conditions reviewed by the responsible actors.

Unplanned failure has a recovery process. Diagnosis, access, spares, repair, testing and release can all consume time before productive service resumes. A repair-completion record may precede operational readiness because the system still needs inspection, crew or a compatible dispatch opportunity. The incident dictionary should distinguish mechanical restoration from return to eligible service. Otherwise a maintenance report can show recovery while the corridor continues to fail its customer commitment.

Common maintenance windows can improve coordination or concentrate vulnerability. Aligning an upstream shutdown with downstream maintenance may reduce otherwise wasted productive opportunities. It may also remove all alternative capability at once. The choice should be evaluated against service commitments, buffer eligibility and recovery readiness. A single annual multiplier cannot demonstrate

which arrangement is better. The operating plan should describe what remains usable during the window and what evidence confirms that the chain can resume together.

Separate intrinsic readiness from the dispatch consequence

A maintenance team can record an asset ready while an operator cannot deploy it. This can occur because the next path is unavailable, the yard has no room or the cargo remains unreleased. The distinction matters to attribution. Mechanical availability should not be penalised for every system delay, but system service should not count the asset as useful merely because repair is complete. A shared record should preserve both measures and the transition linking them.

The maintenance basis should also identify excluded functions. A loader may be available for one cargo type but unavailable for another configuration. A wagon may be mechanically sound but not compatible with the planned route. A control system may support local operation while its interface remains under test. Binary availability can conceal these conditional states. A useful readiness declaration names the intended duty and the limitations that remain after release.

Spares and technical support are dependencies, not universal protection. A stocked component cannot shorten restoration if the failure requires specialised access, diagnosis or regulatory acceptance that remains unavailable. A service contract may promise response rather than completed repair. The corridor plan should distinguish those milestones. It should retain the scope of each commitment and avoid crediting an unproven restoration duration solely because an agreement exists.

Maintenance information should feed the next operating plan. If a repair changes the permissible envelope or leaves temporary restrictions, dispatch must account for them. A local incident closure should not erase those conditions. Likewise, recurring failures should trigger review of the capacity basis and intervention hypothesis rather than continue to be treated as unrelated exceptions. The annual model is only as useful as the readiness assumptions that the operators continue to support.

Use deterministic availability cases honestly

Illustrative calculation. Assumptions and sensitivity are published with the model.

The lower-availability study changes the rail input to 50.00 percent while preserving the other fictional inputs. Calculated flow falls to 7.00 Mt/year and the operating unit cost becomes 29.14 USD/t. This demonstrates sensitivity to a defined annual usable-capacity assumption. It does not estimate how often rail failure occurs or how long any repair would take.

Illustrative. The canonical engine also reports an availability stress for each fixture. For the base chain that deterministic calculation produces 10.50 Mt/year and 24.76 USD/t. The stress changes the input basis according to the app method while leaving demand and costs otherwise unchanged. It is a scenario result, not a probability-weighted expected outcome, confidence interval or forecast. The paper preserves that status whenever the values are displayed.

The distinction prevents a common misuse of stress analysis. A severe scenario can expose a service dependency without establishing its frequency. A mild scenario can leave a ranking unchanged without demonstrating robustness to every relevant interruption. The team should choose additional cases around the mechanisms that could reverse the decision, such as a closure concentrated in a vessel window or a shared loss of power. Those chronological mechanisms lie beyond the annual app and require a suitable separate method.

The cost consequence is likewise restricted. The illustrated unit cost rises because fixed operating obligations are spread over less common flow. Actual failure could also change repair, overtime, storage, contractual and customer costs, none of which is automatically represented. An incident appraisal should not use the annual denominator effect as a complete estimate of loss. It is an explanatory component that should be combined with verified incremental consequences when making a real decision.

Plan restoration as a corridor state

Restoration should be coordinated across actors. An upstream stage restarting at its maximum rate may overwhelm a depleted downstream workforce or constrained yard. A downstream stage seeking to clear a backlog may exhaust the eligible stock needed for the next commitment. The recovery plan should therefore specify the sequence, acceptance checks and temporary priorities that allow the chain to return to a stable rhythm. Local maximum output is not a sufficient recovery objective.

The restoration sequence should distinguish a technical release from a dispatch reservation. In a hypothetical repair, a locomotive is fit for duty after its test, but the path needed to return it to service has been allocated to backlog traffic. Declaring the repair late misattributes a scheduling conflict; declaring the corridor restored ignores the absent assignment. The restoration record should place the technical release, crew readiness, compatible formation and next usable path in sequence. Where the path is the limiting transition, compare protecting a restoration slot with waiting for the normal timetable. This is a calendar choice within legitimate access rules, not a reason for technical staff to waive an operational prerequisite.

Recovery is complete only when the intended service and protective reserves are restored, subject to the agreed definition. Closing an incident after the first successful movement can leave backlog, depleted buffers or unresolved temporary restrictions. A service record should distinguish resumed operation, stable repeating service and restored readiness for the next interruption. These are qualitative milestones here; no unsupported duration or success probability is assigned.

Maintenance alternatives should be compared through displaced service opportunities as well as lost equipment hours. A hypothetical simultaneous closure can avoid repeatedly interrupting the same cargo stream, yet remove the only capability that could recover it during the common window. Staggered work preserves partial capability but may prolong interference with dispatch. The discriminating record is a calendar overlay of the required cargo releases, planned work, usable paths and reserve duties. Mark which commitments remain feasible under each arrangement and which are merely deferred. This reveals a trade-off that equal annual availability can conceal. The lower-availability fixture establishes annual sensitivity; it cannot select between these calendars or price the support and transition burdens they create.

Exhibit: restoration milestones are different

Original maintenance and service architecture. The sequence is conceptual and carries no assumed incident duration.

Exhibit: restoration milestones are different			
Milestone	Evidence	Authority	What remains unresolved
Technical repair complete	Repair and inspection record	Maintenance function	Operational resources and access
Fit for declared duty	Tests and restrictions recorded	Responsible technical authority	Cargo and downstream readiness
Operational deployment ready	Crew, path, cargo and interface confirmations	Operator and relevant interface owners	Backlog and reserve restoration
Stable service resumed	Repeated accepted handoffs under temporary plan	Corridor coordinator within agreed powers	Residual restrictions and future readiness
Recovery closed	Normal service basis and reserves restored	Agreed incident governance	Longer-term corrective action may remain

Maintenance and recovery review

Review the calendar of lost opportunity and the path back to useful service, not just an annual availability percentage. Reconcile planned work, failure restoration and system deployment. Test common dependencies and recovery surges. Require explicit restrictions and accountable releases before treating an asset as ready for the promised service.

1. Record the duty and calendar of each maintenance state.
2. Separate repaired, fit, deployable and service-restored milestones.
3. Coordinate backlog clearance with adjacent capacity and buffers.
4. Update the annual capacity basis when recurring evidence contradicts it.

Maintenance can be coordinated without a detailed corridor simulation

Experienced operators may manage windows effectively through an agreed calendar and interface plan. A simulation is not mandatory merely because timing matters. The evidence should show that the declared service and recovery duties fit the plan and that critical conflicts are understood. Where the choice depends on uncertain concentrated disruptions or finite-buffer behaviour, a more explicit time-based model may be warranted. Model complexity should follow the decision, not replace practical operating knowledge.

The stress is not a failure model

Availability inputs and deterministic stress cases do not establish failure distributions, restoration times or the reliability of a maintenance plan. No event frequency is asserted. The canonical outputs represent annual sensitivity only and require operational evidence before actual service commitments are made.

Evidence and scope

The World Bank lifecycle-cost principle connects construction with operations and maintenance over an infrastructure asset's life. It supports reviewing maintenance obligations alongside initial asset choices, while supplying no repair duration, supplier cost or service warranty for this study.

Evidence: Raising economic efficiency in view of life-cycle costs

Decision implications

- Place maintenance on the relevant operating calendar.
- Distinguish repair completion from useful service restoration.
- Use availability stress as sensitivity rather than event probability.

MECHANISMS AND WORKED STUDIES

12

Weather and seasonality require operating envelopes

An operating envelope and escalation protocol

An operating envelope states when movement is permitted

Weather and seasonality should enter the study through a defined operating envelope. The relevant question is which physical conditions permit the declared service and which conditions require restriction, suspension or a different plan. A generic claim that weather affects logistics is too broad to select an intervention. The team should identify the specific mechanism, such as marine access, road condition, material handling or safe equipment operation, and obtain the applicable technical and jurisdictional basis.

An envelope should distinguish observations, thresholds and actions. An observation records a condition. A threshold is an approved decision boundary for a particular duty. An action states what the operator must do when that boundary is approached or crossed. These elements should not be invented by the analyst to complete a diagram. The responsible technical and operating actors should establish the threshold and its evidence. The corridor manager then coordinates the consequences across adjacent stages.

Seasonality can affect the alignment of opportunities without changing the nominal equipment rate. An inland movement may be restricted during a period when the port remains usable. A vessel window may be lost while the mine continues to release cargo. A storage or dispatch plan can partially bridge that mismatch, but its usefulness depends on eligible stock and recovery. An annual capacity factor can screen the consequence of lost opportunity while concealing how the remaining service must be scheduled.

The study deliberately publishes no climate frequency or local weather threshold. It has no defined geographic empirical universe and no verified operating time series. Assigning a plausible seasonal closure rate would create unsupported evidence. The useful contribution is a method for declaring envelopes and escalation, plus a record of what data would be needed for a quantitative assessment. This preserves the distinction between a physical mechanism that merits attention and an estimated likelihood that has not been established.

Avoid importing regional statistics into a local guarantee

A regional climate or transport statistic can inform context without establishing the conditions at a particular route or berth. The observation location, instrument, time period and operational relevance matter. A rainfall series is not automatically a road-access series, and wind measurements are not automatically a loader-operability record. The analyst should connect the environmental observation to the actual stop rule and the service consequence before using it as a capacity input.

Historical closure records require a consistent operating policy. A change in equipment, access procedure or threshold can alter the relationship between the same physical condition and service interruption. Comparing closure rates across policy editions without a bridge can suggest a trend that reflects changed decision rules. The record should retain both the environmental event and the rule applied. That permits the team to distinguish a change in exposure from a change in capability or tolerance.

Missing conditions can be important. A local record may contain productive periods more reliably than suspended operations, or a sensor may fail during the most severe event. Treating missing observations as benign can bias the capacity basis. The data plan should document sensor coverage and incident reporting, including conditions that were not measured. A quantitative estimate should state how those gaps affect the proposed decision rather than hide them inside a polished annual average.

A future operating envelope also requires caution. Infrastructure works or changes in cargo can alter exposure, drainage, handling and access. A historical series can support a method while failing to represent the proposed operation. The sponsor should review the physical mechanism and changed conditions explicitly. This is a reason to update the capacity basis and challenge scenarios, not an invitation to replace evidence with arbitrary forecast percentages.

Coordinate restriction before the chain fills

A downstream restriction can propagate inland through buffers and queues. If the loader stops while trains keep arriving, the yard may lose usable receipt space. The railway can then block the producer. The escalation plan should identify the point at which upstream release must change, who communicates it and which user commitments are affected. Waiting until every local asset reaches its limit can make a manageable interruption into a system-wide recovery problem.

The plan should preserve safety and lawful authority. A corridor coordinator cannot override a technical stop rule or a public restriction in the name of throughput. Its role is to coordinate service changes within the approved envelope and to make unresolved conflicts visible. Commercial urgency can influence prioritisation among feasible movements, but it cannot convert an ineligible operation into an eligible one. Noncompensable conditions should remain explicit in the decision record.

Communication needs a state definition. A watch condition, restricted operation, suspended operation and approved restart are different states with different consequences. The record should distinguish an estimated recovery from an authorised restart. It should name the source of the physical observation, the decision authority and the interface owners receiving the change. A broadcast warning without a shared state vocabulary can leave actors taking incompatible actions based on their own interpretation.

The upstream response should account for cargo eligibility and user commitments. Restricting mine release may protect yard space while affecting a producer's storage or quality. Rerouting may preserve movement while requiring unconfirmed access. Holding a vessel may alter commercial exposure. The coordinator should record the chosen response, alternatives and authority rather than present it as a technically inevitable consequence. The physical stop condition can be nonnegotiable while the service response still contains important choices.

Restart through evidence rather than optimism

Restart requires the conditions needed for the declared duty to be re-established. A weather observation returning within a nominal threshold may be insufficient if damage, contamination, displaced material or access restrictions remain. The responsible actors should verify the relevant consequences before resuming service. A checklist should relate each check to the mechanism that caused the stop, avoiding a generic green status that conceals unresolved conditions.

The first resumed movement should be distinguished from stable service. Equipment and paths can be available while the backlog, stock configuration or personnel pattern remains disrupted. A controlled restart may deliberately use a reduced operating envelope. The service record should retain that restriction and state what evidence permits return to normal operation. The annual model does not simulate this ramp, so it cannot establish the quantity delivered during the recovery period.

A resilience intervention should be tested against the specific restriction it addresses. Additional storage may protect the loader from inland access loss but fail when the stockyard itself is exposed. Alternate marine access may share the same governing condition. An equipment enclosure may protect one handling function while leaving the approach unavailable. The challenge should trace common causes and deployment requirements instead of counting nominal alternatives as independent protection.

An envelope review should compare observed conditions with the rule in force at the time. If a hypothetical terminal reports fewer weather closures after widening its operating rule, the lower closure count does not by itself demonstrate a milder environment or a more capable asset. Reclassify the old and new observations under a common rule where the records permit it, and retain consequences that the wider rule may expose. If the observations cannot be reconciled, separate the periods rather than infer an improvement. This rule bridge gives the committee a specific test of a resilience claim and identifies whether the proposed change concerns environmental exposure, technical protection or operating permission. It supplies no local forecast or authorised threshold.

Completed envelope and escalation record

The fictional plan names states and authorities while deliberately leaving local physical thresholds unresolved. No climate frequency or operating limit is invented.

Completed envelope and escalation record			
State	Decision basis	Required action	Unresolved evidence
Watch	Relevant environmental condition approaches an approved boundary	Confirm observation and notify adjacent stages	Local sensor and threshold basis absent
Restricted	Responsible actor authorises a narrower duty	Align release, receipt and service commitments	Permitted configurations need technical review
Suspended	Applicable stop condition prevents the declared duty	Protect cargo, reserve space and revise dispatch	Stop authority and local rules not supplied
Restart review	Physical condition improves and consequences are inspected	Verify access, equipment and eligibility	Inspection and authorisation evidence absent
Normal service restored	Repeated handoffs meet the agreed envelope	Close temporary restrictions and restore reserves	No recovery chronology modelled

Operating envelope approval

Do not publish a weather-adjusted service promise until the physical mechanism, local threshold and decision authority are defined. Coordinate upstream actions before storage and queues fill. Keep forecast recovery separate from approved restart. Require an evidence bridge whenever the operating rule or asset configuration changes.

1. Connect environmental observation to the relevant duty.
2. Obtain the threshold and authority from responsible actors.
3. Define watch, restriction, suspension and restart states.
4. Test common causes and the restoration of protected service.

A concept design may need a provisional envelope

A provisional envelope can be useful when it is explicitly a design assumption and when the unresolved technical evidence is scheduled before an irreversible commitment. It should not be presented as an established local threshold or operational permission. The committee can compare concepts conditionally, then require validation of the mechanism and authority before procurement, commissioning or service guarantees. The provisional status must survive later reuse of the model.

No environmental probability is asserted

The study defines no geographic climate universe and reproduces no weather time series. The framework cannot estimate closure frequency, expected loss or future operability. Quantitative claims require relevant verified observations, declared operating rules and a suitable implemented analysis.

Decision implications

- Use local mechanisms and approved thresholds to define the envelope.
- Coordinate restriction across the chain before queues and buffers become exhausted.
- Keep provisional assumptions distinct from operating authority.

GOVERNANCE AND APPLICATION

13 Border processes are part of the production system

A process dependency map with named owners

A physical movement can remain administratively ineligible

Border processes belong inside the corridor service because they can determine whether cargo is permitted to advance. Documentation, inspection, security, health and fiscal requirements can involve several authorities and private actors. The physical availability of a road, rail line or berth does not remove those conditions. A capacity analysis that excludes them can be mathematically correct for unrestricted movement and misleading for the accepted service the user actually needs.

The administrative chain should be mapped by decisions and evidence, not merely by office names. Each step has an applicant, required information, review authority, release condition and downstream consequence. Some steps can proceed before cargo arrives; others depend

on physical inspection or a completed prior decision. The sequence matters because parallelisable preparation differs from a mandatory dependency. The analyst should obtain the applicable local process rather than assume that the fastest imaginable workflow is legally or operationally available.

Cargo status should distinguish submission, completeness, review, approval and physical release. A submitted document may be incomplete; an approved document may not cover the current consignment; an issued release may require another agency's condition before movement. Treating submission as clearance creates false readiness. The operating dictionary should identify the actual event that permits the next handoff and the authority responsible for confirming it.

This paper does not provide a legal opinion about any shipment or jurisdiction. Public primary guidance supports the architecture of pre-arrival preparation, coordinated border processes and release measurement. Actual applicability and implementation remain matters for the relevant authorities and qualified review. The service model should preserve that limit. A generic institutional principle is useful for asking the right question, but it cannot establish an entitlement to move an unnamed cargo across a particular border.

Document quality is an upstream operating resource

A missing or inconsistent document can interrupt the service after equipment and cargo are already committed. The preparatory process should therefore be treated as an operating dependency with its own readiness states. Cargo identity, quantity, classification and destination information should be consistent across the records used by the relevant actors. A data mismatch can create rework even when each participant's local system appears complete.

The process should define who owns the authoritative value for each required field and who can correct it. A producer may know the cargo specification while a carrier controls movement identifiers and an agent submits the release request. The team should avoid forcing an actor to certify information it cannot verify. A shared data dictionary and correction route can reduce ambiguity, but they do not eliminate the need for accountable declarations under the applicable process.

Pre-arrival preparation can move work earlier, provided the relevant procedure accepts it and the information is sufficiently stable. An early filing that is repeatedly revised can create a different burden rather than

remove work. The planner should distinguish information known at dispatch from information available only after weighing, sampling or arrival. The operating benefit comes from feasible sequencing and completeness, not from labelling all administrative work digital or early.

An eligibility gate before irreversible dispatch can prevent avoidable remote holding. The gate should identify the release evidence needed at that point and what may legitimately remain provisional. It should also record the cost of holding cargo upstream. The sponsor should compare alternatives on the whole service, since an upstream hold can reduce border congestion while affecting producer storage or delivery timing. Administrative coordination is a service design choice with operational consequences, not a free capacity improvement.

Inspection and release need a shared observation boundary

A release-time measure should state its start and end events. Arrival at a border area, entry into a controlled facility, document submission and presentation for inspection are different starts. Documentary approval, permission to leave and actual physical exit are different ends. Each can serve a legitimate diagnostic purpose. A performance comparison becomes misleading when one edition uses a narrower interval without explaining the change.

The observation universe also matters. Selected inspections, routine releases and exceptional cases may have different processes. A completed-only average can omit cargo still held at the reporting cutoff. A facility record can omit the queue on the approach. The collection plan should preserve case classes and unfinished states instead of using an aggregate to imply uniform service. These requirements mirror the corridor measurement dictionary because the administrative chain is part of the same accepted movement.

Cause attribution should remain careful. A long release interval can involve incomplete submission, multiple reviews, physical inspection, payment, security or a carrier's own scheduling delay after permission is granted. A report should not attribute the entire interval to a single agency without evidence. The useful output is a sequence of states and accountable decisions, with unresolved causes marked explicitly. This allows coordination without turning the model into an unsupported allegation.

An administrative intervention should be evaluated on release and physical movement together. A process can issue approval sooner while the cargo remains blocked by terminal capacity or transport readiness. That may still improve the user's preparation and certainty, but it is not automatically shorter end-to-end delivery. The committee should request the precise outcome and compare it against the declared service boundary. Better processing statistics and better corridor service are related propositions that require separate evidence.

A release correction should preserve the cargo lineage

A fictional consignment can arrive with an inconsistency between its movement identifier and the declared cargo record. The physical cargo may be intact and the carrier may have completed transport correctly, while the next permission cannot be confirmed. The correction process needs an authoritative relationship between the two objects. Replacing one identifier casually can break the audit trail, create duplicate receipt or leave another record apparently unfinished. The service plan should define who can correct the relationship and how the original remains traceable.

The record should distinguish a clerical correction from a change in the cargo claim. Correcting an identifier is different from revising quantity, grade, destination or classification. Those changes can require further evidence and review under the applicable process. The coordinator should not treat all of them as an administrative edit merely because a digital field can be changed. The next physical handoff should rely on the appropriately confirmed state rather than the convenience of a revised screen.

During the hold, the operating response should protect the cargo and the surrounding service. The yard may need to segregate it, release a reservation or revise the loading plan for other eligible stock. Each choice has an authority and consequence. Additional handling capacity may reduce the burden of segregation while leaving the permission unresolved. The distinction prevents the committee from attributing the entire hold to a physical capacity deficit and helps it identify the specific information or governance intervention.

Closure should retain the original inconsistency, correction evidence, approving actor and resulting release. That record supports later process improvement without asserting blame beyond the evidence. Repeated inconsistencies can justify a dictionary or validation change, but the

evaluation should test whether the new process actually reduces unresolved holds under a comparable basis. The completed correction is a local event; its wider service benefit remains a separate proposition to be established.

Govern the interface without erasing public duties

The corridor governance arrangement should include a route for communicating administrative status to operators within the applicable information-sharing rules. It should identify who can confirm readiness, correct an error and escalate a stalled dependency. It should not assume that private commercial actors are entitled to every authority-held dataset. Data access and confidentiality conditions need explicit review before a shared platform is designed.

Coordination does not mean removing legitimate controls. Health, security, revenue and safety decisions can be noncompensable conditions. The service plan should seek completeness, sequencing and consistent information rather than represent every inspection as inefficiency. An operating proposal that depends on waiving a duty should be rejected or reviewed through the appropriate authorised process. Increased throughput cannot justify ignoring a condition that determines whether the movement is permitted.

A correction request can be routed by the kind of discrepancy. In the fictional identifier mismatch, a transcription correction supported by the original cargo record belongs with the actor authorised to amend that record. A change to the declared cargo, quantity or consignee may require a different review rather than the same clerical route. If the authoritative record is consistent but the receiving system has not ingested the amendment, the remaining intervention is transmission and acknowledgement. Keeping these branches separate prevents a coordinator from repeatedly escalating a system-refresh problem as a substantive release dispute. It also prevents a material cargo change from passing as a harmless identifier edit.

Measure the result of the correction at both the release event and the physical handoff. A hypothetical amendment can clear administrative review while the formation still misses its receiving window. That outcome supports successful document repair and failed service recovery at the same time. Conversely, cargo can move under a permitted temporary arrangement while the permanent record remains unresolved; physical progress then does not establish administrative completion. The

process comparison should preserve both dispositions and the applicable authority. This paired account identifies whether the next action is record completion, transmission, dispatch replanning or destination receipt. An annual administrative capacity input cannot distinguish those paths and should not be treated as a model of individual clearances.

Exhibit: documentary progression and physical progression

Original process architecture for a fictional consignment. The required local procedure and authority have not been established.

Exhibit: documentary progression and physical progression			
State	Evidence that changes it	Actor	Downstream implication
Submitted	Declared information lodged through the applicable process	Applicant or agent	Submission alone does not permit movement
Complete	Required fields and supporting records reconciled	Responsible submitter and reviewer	Review can proceed on a defined case
Reviewed	Relevant checks and decisions recorded	Competent authority	Other conditions may still remain
Released	Applicable permission for the next movement confirmed	Authorised releasing actor	Operator can test physical readiness
Physically transferred	Cargo crosses the declared boundary	Carrier and receiving actor	Service ledger records the actual handoff

Administrative readiness gate

Map the applicable local process and identify which evidence permits each physical handoff. Test feasible pre-arrival preparation and the correction route for inconsistent information. Preserve public duties and information-access limits. Evaluate an intervention on the same release and delivery boundary, including unfinished or excluded cases.

- 1. Separate submission, completeness, review and release.
- 2. Name the authoritative owner of required information.
- 3. Record the actual release event and physical exit separately.
- 4. Escalate with evidence and appropriate authority rather than an assumed right to override.

A private corridor cannot control every border delay

That is precisely why the dependency should remain visible. The sponsor can condition its service promise, improve its own preparation or seek a lawful coordination arrangement. It cannot make the dependency disappear by omitting it from the capacity diagram. The record should distinguish the private actor's controllable actions from external decisions and should avoid assigning obligations to an authority that has not accepted them.

Guidance is not local legal applicability

No source cited here establishes which border procedure applies to a particular consignment, what release is guaranteed or which actor may access protected information. An actual implementation needs jurisdiction-specific review and relevant authority confirmation.

Evidence and scope

The WTO trade facilitation text addresses pre-arrival processing and cooperation between border agencies. Its obligations and implementation depend on member commitments and the relevant jurisdiction. This paper uses the concepts to structure readiness questions and asserts no local entitlement.

Evidence: [Agreement on Trade Facilitation](#)

Evidence and scope

IMO's facilitation overview explains that ship and cargo controls address several public purposes, including health, security and revenue. This supports preserving legitimate control duties within the service map, rather than treating every administrative step as removable delay.

Evidence: [Facilitation](#)

Decision implications

- Track the permission chain alongside the physical chain.
- Treat document completeness as an upstream readiness condition.
- Coordinate within legitimate authority and data-access limits.

GOVERNANCE AND APPLICATION

14

Performance indicators can conceal delivered service

A corridor scorecard with basis definitions

A scorecard should answer a service question

A corridor scorecard should make a decision easier to challenge. It should show whether the declared user service is being delivered, which mechanism threatens it and what action the evidence supports. A page of local utilisation, gross tonnage and average dwell can look comprehensive while failing to answer those questions. The designer should begin with the service boundary and a small set of diagnostic relationships, then retain the underlying definitions so that the reader can reconstruct each measure.

Delivered quantity needs an acceptance basis. Gross handling can include recirculation, rejected cargo or movement before final release. A local transfer may count progress that has not reached the customer boundary. The headline should state the quantity basis and event that enters its numerator. A separate local indicator can remain useful for operations, provided it is labelled as such. The scorecard should not force all actors to use one number when their legitimate purposes differ, but it should explain how their numbers reconcile.

Timeliness should describe the commitment and population. An on-time measure needs a defined due event and treatment of revised promises, cancellations and unfinished movements. If the schedule is repeatedly moved after a delay, apparent punctuality can improve without the user receiving a better service. Preserve original and revised commitments so that the report can distinguish planning adaptation from fulfilment. A punctuality claim without that basis invites a favourable interpretation that the evidence may not support.

A diagnostic indicator should also have an owner and response. A growing eligible queue can trigger dispatch coordination; missing release information can trigger a documentary action; loss of reserve can trigger a recovery plan. A metric with no decision consequence is at risk of becoming reporting furniture. The committee should ask what the organisation will do when the indicator changes and who can legitimately make that choice.

Averages hide distribution and unfinished service

An average duration can be useful for a defined population, but it does not describe every user's experience. A small group of severe delays can matter commercially even when the central tendency changes little. Conversely, a change in cargo mix can alter the average without any process becoming worse. The report should retain meaningful case classes and show the evidence relevant to the service commitment. It should avoid publishing tail measures that have not been supported by a suitable observation universe.

Unfinished service should remain visible at the reporting cutoff. Completed movements tell one story; cargo still waiting tells another. A scorecard that reports only completed cases can improve when the most difficult work remains unresolved. The dashboard should show the unfinished population and its states, with durations treated as incomplete. It should also preserve cancellations and diversion where they matter to the user outcome rather than quietly remove them from the performance universe.

Period comparisons require an edition bridge. A new timestamp, route boundary, cargo eligibility rule or sensor coverage can change a metric without changing the underlying service. The report should record that change and decide whether the periods remain comparable. A trend line that joins incompatible definitions can create a false performance story. Where a bridge is not feasible, show a break and explain the new basis instead of prioritising visual continuity over evidential continuity.

The scorecard should separate capability from use. Low use of a spare stage can reflect a binding neighbour, demand limitation or a deliberate recovery reserve. High use can reflect efficient service or a fragile absence of slack. Neither percentage interprets itself. The annual app's utilisation is a ratio within its fictional capacity balance, while an operational scorecard may use productive time, duty availability or resource occupancy. The labels and definitions should make that distinction clear before comparisons are made.

Do not confuse logistics indices with site performance

A broad logistics index and a corridor operating measure answer different questions. A perception-based index can describe respondents' assessments of a national logistics environment under its methodology. A site record can describe movements within a declared facility or route universe. Neither should be substituted for the other. A favourable

national context does not establish a particular stage's capacity, and a local delay episode does not determine an entire country's logistics performance.

Methodology changes further limit comparison. The source cards retained for the historical and redesigned World Bank logistics indicators describe different evidence bases. The paper does not treat them as an unchanged time series. It also imports no country score, ranking or site capacity. Their useful role here is to demonstrate why a scorecard needs an evidence-object definition and an edition bridge. A new indicator can be valuable precisely because it measures something different.

A corridor committee can use external context alongside internal evidence if the roles are clear. Context can suggest questions about institutional dependencies or service uncertainty. Internal records can test the mechanism within the actual service boundary. The decision should state which source supports which proposition. An external label should not supply an empirical parameter merely because the project team lacks local observations, nor should a local model result be presented as validation of the wider index.

The same caution applies to benchmarking across ports or routes. Cargo mix, boundaries, timestamps, operating conditions and acceptance rules can differ. A comparison can still be informative if those differences are declared and the decision consequence is appropriately narrow. The team should resist a single ranking that conceals incompatible bases. An original comparison table showing scope and limitations can be more useful than a league table whose apparent precision exceeds its evidence.

Build a scorecard with paired measures and challenges

The scorecard can pair a service outcome with the mechanism that explains it. Accepted delivery should be accompanied by eligible demand and unfinished commitments. End-to-end duration should be accompanied by state intervals and missingness. Reserve readiness should be accompanied by the protected interruption and restoration plan. This structure gives the reader both the result and the conditions needed to interpret it, reducing the temptation to treat a favourable local measure as a complete system judgement.

Test a scorecard revision with a hypothetical queue moved from the gate to an appointment hold upstream. The terminal's gate-to-receipt measure improves because its start event now occurs later in the user's journey. The end-to-end accepted-delivery measure may remain unchanged. Add the

held consignments to the same service cohort and retain the original planned arrival, its revisions and actual receipt. If the local indicator still improves while the cohort indicator does not, report an internal process improvement with displaced waiting, rather than a total service gain. This adversarial comparison makes the measure's start event operationally consequential and exposes the effect of changing where eligible demand becomes visible.

A challenge column can identify what would make the current interpretation wrong. If low delivered flow is attributed to rail capability, the scorecard should ask whether demand, cargo eligibility or destination receipt explains it instead. If a shorter delay is presented as an intervention benefit, identify the concurrent changes that could account for it. The challenge need not paralyse action. It directs the next observation or trial and prevents a dashboard from becoming a one-way presentation of management's preferred explanation.

A scorecard can carry a reconciliation disposition instead of forcing competing reports into a single total. When production release rises while accepted delivery is flat, reconcile the difference into in-transit eligible cargo, unfinished receipt, rejection, loss or unresolved identity. When both quantities are flat but completion time improves, retain that time benefit separately from throughput. When a denominator shrinks because users withdrew, show withdrawn demand rather than celebrating utilisation. These branches distinguish a service change from a reporting change and locate the next investigation. The fictional capacity balance supplies no observations for this reconciliation; its contribution is the structural bound against which an actual cohort account can be compared.

Exhibit: outcome and diagnostic basis

Original scorecard architecture. Values remain unpopulated because the study contains no actual operating universe.

Exhibit: outcome and diagnostic basis			
Measure	Declared object and boundary	Companion evidence	Interpretation challenge
Accepted delivery	Eligible cargo accepted at the chosen destination	Demand, rejection and unfinished commitments	Does gross handling overstate useful service?
End-to-end duration	Defined consignment from eligible release to acceptance	State intervals, unfinished cases and missing clocks	Was waiting moved outside the reported boundary?
Annual capacity balance	Canonical fictional serial chain	Nominal rates, availability and demand assumptions	Can chronology or quality make the bound unattainable?
Reserve readiness	Deployable protection for named interruption	Eligibility, access and restoration record	Does common cause disable the reserve?
External logistics context	Source-specific population and methodology	Edition and jurisdiction bridge	Is context being misused as local capacity evidence?

Scorecard publication gate

Publish a measure only with a defined service question and reconstructable basis. Show unfinished service and consequential exclusions. Pair outcomes with diagnostic evidence and an accountable response. Mark a methodology break when definitions change, and prevent external indices from becoming unsupported local parameters.

- 1. Define object, numerator, denominator and acceptance event.
- 2. Preserve original commitments and changes to them.
- 3. Show missingness and unfinished cases alongside completed results.
- 4. State the alternative explanation that the next review should test.

Senior readers need a concise dashboard

Conciseness is compatible with rigorous definitions. The main page can show the few measures that affect the decision, with a compact basis reference and an explicit unknown where necessary. It should not compress several incompatible service objects into a single unexplained score. A senior reader benefits from knowing which uncertainty can reverse the recommendation more than from receiving a dense panel of favourable numbers.

No operational ranking is produced

The paper has no measured corridor universe and imports no external index score as a stage capacity. Its scorecard is a proposed architecture. A real ranking or performance claim requires compatible evidence, defined periods and a bridge for methodology differences.

Evidence and scope

The historical World Bank report describes survey-based logistics assessment alongside complementary operational measures. Its methodology is not a site-capacity measurement and supplies no empirical input to the fictional chain.

Evidence: *Connecting to Compete 2023*

Evidence and scope

The World Bank's redesigned logistics indicators use operational shipment data, differing from the survey basis of earlier editions. This change requires a methodology bridge rather than an assumed unchanged series.

Evidence: *About Logistics Performance Indicators 2.0*

Decision implications

- Design measures around accepted service and action.
- Keep unfinished movements and definition changes visible.
- Separate external context from local operating evidence.

GOVERNANCE AND APPLICATION

15 Commercial incentives can create queues

An incentive-incidence matrix

A local incentive can contradict system service

Actors can behave rationally under their own agreements while worsening the corridor outcome. A terminal rewarded for productive equipment time may prefer large ready batches, leaving smaller urgent commitments waiting. A producer rewarded for evacuation may dispatch cargo before downstream eligibility is resolved. A carrier rewarded for departure may accept a movement that cannot be received promptly. The resulting queues are not necessarily evidence of incompetence. They can reflect the objectives and costs assigned to each actor.

The incentive map should identify who gains, who bears cost and who can change the relevant decision. A storage charge may discourage excessive inventory or make prolonged occupancy commercially attractive to another party. A priority fee may support scarce capacity management

or transfer delay to ordinary users. A departure target may improve local activity while creating destination congestion. The map should describe the mechanism without assuming motive or alleging misconduct unsupported by evidence.

A commercial term has a defined legal scope. Its meaning, enforceability and exceptions cannot be inferred from a general operating model. This paper uses no private contract or actual tariff. It proposes an analytical structure for reviewing how agreed incentives may interact with service. An actual implementation would require the relevant agreements and qualified review. The distinction matters because an apparently perverse outcome can be consistent with a legitimate risk allocation that the sponsor intentionally accepted.

The service objective should therefore sit alongside local incentives before operations begin. If the corridor needs recovery reserve, an agreement that rewards only immediate utilisation may make that reserve difficult to sustain. If the corridor needs early information, a process that penalises every forecast revision may discourage honest updates. The governance design should recognise the response that the rule creates rather than assume that all participants will prioritise the system without an agreed basis.

Charges can reveal scarcity or entrench delay

A charge for occupying a scarce resource can encourage users to release it. The effect depends on whether the charged actor controls the release. If a consignee pays storage while awaiting another authority's decision, the charge may not change the bottleneck. If the terminal cannot accept the next cargo because reserved stock remains, a release incentive may help. The analyst should trace the charge to a controllable action rather than judge it solely by its name or apparent severity.

A delay-related charge can also redistribute risk. A vessel, terminal and cargo owner may have different commitments and different exposure when service is late. Their agreed terms determine who bears the financial consequence, while the physical cause can sit elsewhere in the corridor. The incentive review should preserve those distinctions. A payment obligation does not establish technical causation, and a technical constraint does not decide legal liability.

Local revenue can depend on states that the corridor wants to minimise. This creates a potential conflict that should be examined through actual terms and evidence rather than assumed. The governance arrangement can align service objectives, share information or change the basis of

remuneration where the parties agree. It should not presume that removing a charge will improve capacity. The charge may fund readiness or discipline scarce use; its replacement needs a coherent operating and commercial basis.

Priority mechanisms deserve particular scrutiny because they move delay between users. A paid or contractual priority can be legitimate, but the service impact should be visible. Repeated emergency priority can undermine the normal plan and make forecasts less dependable. The record should distinguish agreed service classes from discretionary exceptions. The committee should know whether the corridor's claimed performance represents all users or only the class that receives preferential access.

Information incentives matter as much as asset incentives

The corridor needs accurate readiness information before every condition is final. An actor who is penalised for uncertain forecasts may delay reporting until certainty arrives, leaving little time for coordination. An actor who is rewarded for optimistic schedules may repeatedly announce readiness that does not occur. Both behaviours can damage service. The information agreement should distinguish a forecast, a commitment and an actual event, with appropriate accountability for each.

Data sharing can also impose costs or expose commercial information. A private actor may resist a platform that requires detailed operational records without explaining their purpose, protection or access rules. The sponsor should define the minimum information needed to coordinate the service and the legitimate recipients. A governance charter that assumes unrestricted data access is likely to fail or create unacceptable disclosure. No external transmission of private operating data is proposed by this paper.

The escalation process should support candour. A recurring delay reported promptly can lead to useful intervention; an adverse local score can lead an actor to recode the state or narrow its measurement boundary. The scorecard should therefore pair accountability with a way to challenge definitions and explain circumstances. This does not excuse poor performance. It makes the evidence more reliable by preventing the reporting system from rewarding concealment of the mechanism the corridor needs to understand.

An intervention can change who bears uncertainty. Earlier cargo release may make the producer carry more stock. A stricter appointment system may transfer waiting to the carrier. A faster recovery plan may require overtime or reserve readiness from another actor. The decision record should identify those transfers and the agreed compensation or service basis where relevant. A system efficiency claim that omits incidence can be technically plausible and commercially unworkable.

Paying for reserve is a service-design choice

A fictional operator is asked to keep a compatible resource ready for recovery while its normal remuneration rewards productive use. The request creates an explicit economic tension. Idle readiness has cost, and assigning the resource elsewhere can be locally rational. The sponsor should first define the reserve duty and deployment condition, then decide whether the existing agreement supports it. A vague instruction to cooperate leaves the operator carrying an unpriced obligation and the corridor relying on reserve that may disappear.

The arrangement could define a readiness service, a deployment service and a restoration duty separately. Each would need evidence and agreed terms. Readiness requires compatibility, personnel and access; deployment requires authority and a feasible operating path; restoration requires return to the protected reserve state. The paper specifies no payment amount or legal form. It identifies the service objects that an actual commercial review would need to consider.

The committee should also challenge whether reserve is the right remedy. If the resource shares the failed infrastructure or cannot reach the relevant duty, paying for readiness does not protect the service. If better release information removes the recurring disruption, reserve may be unnecessary for that mechanism. Conversely, a genuinely deployable resource can be valuable even without increasing normal annual flow. The incentive review should therefore follow the resilience test rather than assume that a commercial payment makes the reserve technically effective.

Evaluation should preserve both operational proof and incidence. The sponsor should know whether the paid duty remained available, whether it could be deployed and which user commitments were protected. It should also know which normal uses were displaced and which restoration costs occurred. These questions make the agreement review concrete without turning an annual capacity balance into an unsupported contract valuation.

Align incentives through a concrete service agreement

The incentive review should begin with a recurring decision that affects the service. Examples include releasing cargo into a full yard, reserving a vessel window, granting exceptional priority or holding reserve equipment ready. For each decision, identify the actor's local objective, the corridor consequence, the controllable action and the proposed alignment. This is more useful than a broad statement that stakeholders need cooperation, because it shows exactly where a rule changes behaviour.

A revised arrangement should be testable. If a readiness-reporting change is intended to improve planning, measure the notice available and the resulting missed handoffs under a defined basis. If a reserve payment is intended to support recovery, verify deployability during the named interruption. If an appointment rule is intended to reduce queues, include the external approach and upstream holding. The evaluation should preserve the service boundary and distinguish the agreed commercial mechanism from a causal effect that still needs evidence.

The committee should compare feasible alternatives rather than assume that all system benefits can be shared automatically. A rule that improves total accepted service can harm a user whose access obligation must be respected. Some conflicts need negotiation, some need a different service design and some make the proposed intervention inadmissible. The governance arrangement should state how those choices are made and which actors retain authority. A technical capacity model cannot resolve them by assigning an aggregate gain.

Apply the incentive map to a hypothetical forecast that reserves scarce receiving space. If reservation carries no consequence when readiness is repeatedly overstated, an actor can protect its own priority while leaving others without usable slots. A rule that expires unconfirmed reservations changes that incentive, but may penalise genuinely uncertain upstream work. Compare confirmation effort, unused reservations, reallocation and later accepted service under both rules. A permitted readiness payment is a different mechanism: it purchases a defined information or reserve duty rather than penalising uncertainty. The appropriate choice depends on which behaviour creates the conflict. Neither the annual flow result nor a general appeal for cooperation establishes that one commercial arrangement will improve behaviour.

Exhibit: incentive incidence at recurring decisions

Original qualitative map. It describes possible mechanisms and does not assert any actual contract term or actor motive.

Exhibit: incentive incidence at recurring decisions			
Decision	Local incentive	System consequence to test	Alignment question
Dispatch cargo early	Meet evacuation or departure target	Yard congestion and ineligible stock	Is downstream readiness a release condition?
Reserve a ready asset	Protect preferred customer window	Other eligible users lose access	When does the reservation expire or require confirmation?
Hold reserve capability	Avoid unused resource cost	Recovery margin disappears	Who pays for the declared reserve duty?
Report a forecast revision	Avoid apparent local failure	Late notice reduces coordination options	Are forecast, commitment and actual event distinguished?
Apply a delay charge	Allocate cost or encourage release	Charged actor may lack control	Does the rule act on a feasible corrective decision?

Incentive-incidence review

For every recurring decision that can constrain service, identify the actor, objective, cost bearer and controllable action. Test whether the rule supports the declared corridor outcome. Preserve legal scope and access obligations. Evaluate a proposed change through its operating mechanism and the burdens it transfers to other users.

1. Map a specific decision rather than a generic stakeholder interest.
2. Separate physical causation from contractual liability.
3. Identify which actor can change the limiting state.
4. Record the service, information and reserve duties needed for alignment.

Local optimisation can be the efficient allocation

A local rule may correctly reflect scarce resources and agreed risk allocation. The paper does not assume that every divergence from a system target is harmful. It asks whether the target and agreements are compatible and whether the claimed service benefit includes the incidence of the rule. A proposed change should be rejected when it violates a noncompensable obligation or shifts burdens without a workable agreement, even if an aggregate physical measure improves.

No strategic or contractual outcome is estimated

The study contains no actual tariff, private agreement or behavioural dataset. It cannot establish an actor’s intent, legal liability or response to a commercial change. The incidence map is a structured question set for relevant evidence and review.

Evidence and scope

The World Bank risk guidance distinguishes operating, demand, financial and legal mechanisms. These categories help keep physical causes, commercial exposure and authority separate; they provide no event likelihood or contractual conclusion.

Evidence: Identifying Risks

Decision implications

- Trace incentives to controllable recurring decisions.
- Keep physical causation distinct from legal cost allocation.
- Include transferred burdens in any system improvement claim.

GOVERNANCE AND APPLICATION

16

Redundancy is useful only if it is usable

A resilience test for alternate capacity

Alternative capacity must have a deployment path

Redundancy is useful only when it can deliver the protected service under the interruption being considered. An alternate road, spare machine or second contractor can look reassuring on an asset map while sharing the decisive constraint or lacking a deployment right. The resilience review should begin with a named service and failure mechanism. It should then trace the alternative from activation through accepted delivery, including permissions, compatible cargo, personnel and the next-stage handoff.

A route can exist physically without being usable for the relevant load. A spare asset can be mechanically ready without a qualified operator. Another terminal can have nominal capacity without accepting the cargo grade or granting access in the needed period. These are not minor qualifications to an otherwise demonstrated reserve. They determine whether the alternative is inside the feasible response set. The committee should credit only the capacity for which a realistic deployment path has been established, while keeping conditional options visible as unresolved opportunities.

The time required to activate an alternative can govern its value. A reserve that becomes available only after the customer window closes may protect longer-term continuity without protecting that particular delivery. This paper does not calculate activation duration, but the record should distinguish immediate substitute, recoverable alternative and longer-term redesign. Those states support different service promises. Counting them all as equivalent backup obscures which commitments remain exposed during transition.

Alternative capacity also competes with existing users. A second route or terminal may appear spare in normal reports while being contractually committed or needed for other services during disruption. The access basis should therefore include the conditions under which the corridor can actually use it. A general expression of interest is not a firm operating right, and a right is not demonstrated compatibility. The resilience claim should preserve each evidential milestone rather than compress them into a single available-capacity label.

Common causes make apparent alternatives dependent

Two alternatives can share a resource or condition that defeats both. Rail operators can use the same bridge and destination yard. Road alternatives can cross the same vulnerable access point. Spare handling equipment can depend on the same power supply or control system. A port alternative can share the same marine approach or cargo release process. The resilience map should identify those common dependencies before the alternatives are credited as independent protection.

The relevant common cause follows the declared interruption. A spare locomotive can address a traction failure while offering no protection against an unavailable path. An alternate road can address a local route obstruction while offering no protection against a cargo eligibility hold. A stockyard can protect loading from inland interruption while failing during a yard-wide power loss. The analysis should test the mitigation against the actual mechanism, not against a generic loss of capacity that makes every reserve appear helpful.

Concentrated recovery demand can also couple alternatives. After a regional disruption, many users may seek the same spare route or terminal. The corridor's normal access experience may not describe that state. An actual resilience study should obtain the relevant priority and deployment terms and consider shared-demand conditions. The present

study asserts no probability of simultaneous demand. It identifies a mechanism that can make apparently available capacity unusable when it matters most.

Independence should therefore be a claim with evidence, not a visual convention. Separate boxes on a diagram do not establish independent failure or independent access. The map should show shared links, permissions and activation resources. Where the evidence is incomplete, the alternative should remain conditional. This improves the decision because it directs the next diligence action toward the dependency that could invalidate the reserve rather than toward another nominal capacity quotation.

Capacity ties and demand challenge the reserve story

Illustrative calculation. Assumptions and sensitivity are published with the model.

The tied-constraint study calculates 12.00 Mt/year with Rail corridor and Vessel loading jointly binding. This is a deterministic illustration of coupled annual limits. Relieving one equal constraint alone can leave the other unchanged. The app tests single-stage expansion while preserving the rest of the chain, making the lack of system gain visible. The case is not a reliability estimate, but it is a useful challenge to a claim that one added resource necessarily creates useful reserve.

Illustrative. The changed-demand study calculates 8.00 Mt/year and has no binding capacity stage. Every stage has effective annual capacity above the chosen demand. This reverses the expansion interpretation: additional capability cannot create greater common flow while the demand bound remains unchanged. Spare capacity can still have a continuity role, but that role must be defined separately. The scenario prevents the paper from equating more capacity with more service under every admissible condition.

Illustrative. The zero-service study sets vessel-loading availability to 0.00 percent and calculates 0.00 Mt/year. The common annual flow vanishes, and operating unit cost is unavailable rather than a meaningful finite cost per tonne. This boundary case shows that spare upstream capacity does not substitute for an essential unavailable serial stage. It also shows why the service model should treat undefined unit cost honestly instead of presenting a favourable placeholder.

None of these annual cases establishes the value of a real redundant route. They test structural claims about serial flow, tied limits and demand. A true alternative route would change topology, which the app does not

implement. The reader should not add an alternate path as another serial stage and interpret the result as network resilience. The correct next step is a method that represents routing, shared links and access, with inputs and outputs frozen under its own declared model contract.

Accept reserve through an operating proof

The reserve record should state activation, compatibility, access, service duty and return to normal operation. It should identify the evidence for each condition and the remaining gap. A controlled deployment test can verify some interfaces, while a document review can establish a right or restriction. Neither alone establishes the full protected service. The committee should distinguish a reserve identified, a reserve authorised, a reserve tested and a reserve demonstrated under the intended operating envelope.

The test should not create an unsafe or unauthorised disruption merely to demonstrate seriousness. It should select feasible exercises that address the decisive interfaces. A substitution test can confirm equipment compatibility; a dispatch exercise can confirm communication and approvals; a controlled cargo movement can confirm receipt. Residual conditions should remain explicit. A partial exercise can provide useful evidence without justifying a complete continuity guarantee.

The cost of reserve should be compared with the duty it supports. Readiness can require maintenance, trained personnel, access payments or inventory. A nominal alternative that is inexpensive to identify may be costly to keep deployable. Conversely, a spare stage can have little annual throughput value while avoiding a critical service loss. The annual operating model does not calculate these economics. Its role is to prevent spare capacity from being misrepresented as additional common flow without a service mechanism.

A reserve exercise should test the intersection of prerequisites, not merely confirm each in isolation. In a hypothetical backup route, the compatible vehicle is demonstrated on one occasion and the access right is confirmed for another operating window. Those results do not yet establish that the same eligible cargo can use both together during the protected interruption. The reserve ledger should link the demonstration to the cargo, window, receiving condition and shared dependencies that actually coincide. If their intersection is empty, the candidate remains a potential alternative rather than usable protection. This test separates a fragmented collection of favourable proofs from a coherent deployment path without inventing an activation time or failure probability.

Exhibit: usable reserve ledger

Original resilience architecture. The fictional entries are conditional and carry no invented activation duration or probability.

Exhibit: usable reserve ledger			
Alternative	Protected mechanism	Conditions for use	Unresolved common dependency
Spare compatible locomotive	Traction resource unavailable	Authorised crew, compatible formation and path	Same route and destination acceptance
Alternate road	Declared local road obstruction	Loaded-route compatibility and access permission	Shared approach, border or receiving terminal
Released yard stock	Temporary loss of inland replenishment	Correct grade, accessible reclaim and valid release	Yard power, weather and loading permission
Alternate loading facility	Declared local loading failure	Cargo acceptance, marine service and access agreement	Shared approach or documentary release process

Usable redundancy test

Credit an alternative only against a stated interruption and protected service. Trace activation through accepted delivery and identify common dependencies. Preserve conditional rights and incomplete testing. Choose an appropriate network or time-based method when the alternative changes routing or chronology, rather than stretch the annual serial app beyond its scope.

1. Name the failure mechanism and customer duty.
2. Confirm compatibility, access and activation authority.
3. Map shared links and competing demand.
4. Record what the deployment test establishes and what remains untested.

Strategic optionality can justify an untested alternative

An option can be worth acquiring before every operating condition is resolved, especially when access opportunities may disappear. The decision should state that it buys a conditional option rather than demonstrated reserve. The committee should identify the remaining expenditure, evidence and authority needed to make it usable. A capacity diagram should not promote that option into a guaranteed continuity mechanism merely because the strategic rationale is plausible.

No independence or failure probability is calculated

The serial model has no probabilistic dependence or network routing. It cannot estimate the likelihood that alternatives fail together or the expected value of resilience. The framework directs evidence collection and preserves conditional scope.

Evidence and scope

The Orange Book connects risk treatment, objectives, governance and monitoring. Its UK public-sector scope is retained; it supplies no corridor probability.

Evidence: The Orange Book: Management of Risk, Principles and Concepts

Decision implications

- Credit deployable alternatives against named interruptions.
- Show shared dependencies rather than infer independence from separate assets.
- Keep strategic options distinct from demonstrated reserve.

GOVERNANCE AND APPLICATION

17 Expansion should follow a bottleneck intervention order

An intervention prioritisation and stop-go protocol

Intervention order should follow the mechanism

An intervention sequence should begin with the most decision-relevant uncertainty, then compare feasible changes to the mechanism that governs service. The sequence need not always start with physical expansion. A measurement correction can reveal that the supposed bottleneck was a boundary error. A release rule can remove an eligibility hold. A dispatch change can align opportunities. A buffer can protect a temporary mismatch. A permanent asset addition can increase capability when those measures do not resolve the binding duty.

This is a decision order rather than a universal ranking of cheap actions before expensive ones. A physical investment may be urgent or may require a long lead time. An operating change may be inadmissible, costly to users or insufficient for the service. The committee should compare the alternatives on their actual mechanism, feasibility and consequences. The value of the order is that it makes the reason for irreversible commitment visible instead of treating asset addition as the default response to every shortfall.

The first question is whether the reported constraint is real under the approved service basis. If a throughput number mixes gross and accepted cargo, more capacity does not repair the evidence. If observed output is demand-limited, a local equipment upgrade cannot create a customer. If a permission prevents movement, a faster machine cannot override it. These checks should precede a capacity ranking because they determine whether the alternatives address the right problem.

The next question is whether a reversible action can test the mechanism. A focused observation or controlled operating trial can provide evidence before procurement. The trial should have a defined change, comparison basis, stop condition and decision consequence. It is not a delaying tactic when the uncertainty can reverse an expensive choice. It is less useful when the remaining issue is already established and the trial cannot materially improve the decision.

Separate physical dominance from economic preference

The canonical studies establish different physical outcomes under a fixed fictional basis. Loading-only expansion leaves annual common flow unchanged; rail relief reaches loading; coordinated relief reaches the stockyard. These contrasts do not include capital costs, implementation times or incremental operating obligations. A physical ranking is therefore an input to appraisal, not its conclusion. The committee should ask what each alternative costs and what accepted service it changes over the relevant period.

An alternative with greater ultimate capacity may deliver later or create more transition exposure. A smaller intervention may protect near-term service while preserving a future option. A coordinated construction programme may reduce mobilisation costs while increasing the consequence of a common delay. These comparisons require a defined economic and implementation basis. The annual app cannot adjudicate them by displaying a lower operating unit cost under unchanged assumptions.

Feasibility gates should precede compensating comparisons. Safety, legal access, cargo acceptance and mandatory environmental or public obligations may rule an option out. The sponsor should not assign a high benefit score that appears to outweigh an ineligible operating condition. Once feasible alternatives are identified, the committee can compare

physical service, cost, timing, resilience and distributional consequences. The distinction between eligibility and preference keeps the decision logic coherent.

Uncertainty should be connected to the action. If the next constraint is uncertain but cheap to observe, the plan can commission that evidence before committing to a second expansion. If demand uncertainty dominates every physical option, more detailed equipment analysis may add little. If a long-lead component creates a genuine timing constraint, a conditional commitment may be justified while other evidence develops. The record should explain why the selected sequence fits the decision rather than invoke caution or speed as a universal principle.

Use the next constraint as an explicit handover

Every intervention should state what it expects to become binding next. That prediction gives the receiving actor a design and operating brief. If rail relief is expected to make loading binding, the port should review cargo readiness and vessel windows. If coordinated relief is expected to make the yard binding, the yard should examine receipt, segregation and reclaim. The next constraint is not merely a label in a model output. It is a request for evidence about the new service duty.

The prediction also supplies a falsifier. If the intervention does not change the original constraint under a compatible basis, its mechanism may have failed or the input may have been wrong. If a different unexpected constraint appears, the system map should be revised. The sponsor should preserve that learning rather than redefine the project benefit after commissioning. An honest revision can improve the next decision and protect the committee from repeating the same incomplete model.

The annual slate includes cases that challenge the intervention order. Tied constraints can require joint relief; lower demand can make capacity addition irrelevant to flow; unavailable loading can stop the entire serial chain. These are structural counterexamples to a simplistic instruction to enlarge one smallest asset. The planner should identify comparable conditions in the actual service before selecting the intervention. It should also state when the annual method is insufficient because routing or chronology determines the result.

A staged plan needs conditions for continuation and stopping. Evidence of increased local rate alone should not release the next capital stage if accepted service remains unchanged. Conversely, a successful bounded trial can justify further commitment when the mechanism and adjacent readiness are demonstrated. The stop rule should identify which adverse

observation requires revision, suspension or abandonment. A plan with only success milestones can preserve momentum after its original rationale has been contradicted.

Write the committee record around the choice

A committee can use a branching intervention record. For the fictional annual-flow objective, loading-only expansion fails the flow branch while rail relief advances it to the next loading constraint. Coordinated relief advances that branch further but hands receiving duty to the stockyard. Lower demand terminates the expansion branch unless a separate continuity objective is adopted. Write these dispositions beside the retained controlled cases, then attach the actual cost and feasibility work required for any real commitment. This gives the committee an order of questions with explicit exits. It prevents a successful local design from keeping an expansion programme alive after the governing service proposition has disappeared.

A staged authorisation needs a recovery position if the expected transition fails. Suppose a hypothetical operating trial changes dispatch before a physical expansion is committed. If accepted service rises and the predicted next constraint appears, the design brief can incorporate that duty. If local activity rises but accepted service does not, restore the previous feasible dispatch rule and investigate the unexplained handoff. If the trial exposes a safety or access condition, stop that branch and refer it to the responsible authority. These outcomes are different from a vague approval to optimise logistics. They define what the implementing team can do after learning, including the route back to an already workable service.

The review trigger can be tied to the intervention's causal premise. A changed cargo grade reopens compatibility and handling assumptions; a changed destination reopens receipt and cycle return; a changed access window reopens chronology. An unrelated document revision need not reopen every design choice. In the hypothetical rail-relief branch, a new receiving restriction matters because it can consume the additional circulating opportunity that justified the intervention. The trigger record should identify the affected branch and the smallest comparison that can determine whether it survives. This narrows reanalysis to the changed mechanism while preserving the earlier decision as a valid record of its own conditions.

The resulting sequence has a testable termination condition: further expansion stops when the approved service objective is met, demand caps the useful flow, or the next feasible increment cannot justify its full

consequences. A constraint's existence alone is insufficient reason to enlarge it. A corridor always retains some limiting function under a finite design. The committee's task is to decide which remaining limitation matters to the promised service and which belongs in the accepted operating envelope. This distinction converts the bottleneck method from an endless asset programme into a bounded development choice. The annual cases reveal physical transitions; real stopping requires the demand, cost and operating evidence for the actual service.

Exhibit: intervention order and release evidence

Original decision sequence. It is not a universal priority ranking and carries no investment return calculation.

Exhibit: intervention order and release evidence			
Decision stage	Candidate action	Evidence to advance	Reason to stop or revise
Define and verify	Correct service boundary and measurement basis	Coherent object, calendar and acceptance events	Reported shortfall is a definition mismatch
Resolve controllable mechanism	Trial release, appointment or dispatch change	Comparable accepted service and interface records	Waiting is merely displaced or duty violated
Protect bounded mismatch	Test buffer or deployable reserve	Named interruption and feasible restoration	Reserve cannot be used under common cause
Increase capability	Design physical stage intervention	Feasible duty, cost and next-stage readiness	Next constraint absorbs the claimed service gain
Commit coordinated programme	Approve linked interventions and governance	Whole-chain acceptance plan and conditional releases	Demand, access or operating basis no longer supports the choice

Intervention stop-go protocol

Advance only when the proposed action addresses the declared limiting mechanism and its feasible service consequence is supported. Keep investigation, procurement, operation and service acceptance as distinct approvals. State what would falsify the expected effect and who owns the next constraint. Review the choice when the service basis changes.

1. Verify the problem and noncompensable conditions.
2. Compare a reversible mechanism test with irreversible alternatives.
3. Add incremental cost, timing and user consequences to the physical ranking.
4. Release the next stage only on the evidence stated in the approval.

A long-lead project cannot wait for every uncertainty to close

The committee can accept uncertainty deliberately where waiting has material cost. It should identify the remaining assumptions, conditional commitments and options to alter the programme. The key requirement is that uncertainty remains connected to a decision and consequence. A long lead time does not convert an unverified annual input into an established service capability or make the omitted capital economics irrelevant.

This is an appraisal architecture, not an appraisal result

The paper contains no actual investment costs, implementation schedule or customer value. It cannot choose the economic optimum for a project. The intervention order structures the evidence and approvals needed to make that choice.

Evidence and scope

The Green Book separates objectives, options, costs, benefits and uncertainty in UK public appraisal. Its conventions are not automatically private-finance assumptions.

Evidence: [The Green Book](#)

Evidence and scope

GAO's cost guidance emphasises technical baseline, assumptions, sensitivity and updates with actuals. It supplies no intervention cost for the fictional corridor.

Evidence: [Cost Estimating and Assessment Guide](#)

Decision implications

- Sequence actions around the mechanism and value of evidence.
- Add full costs and feasibility before selecting an investment.
- Make stopping conditions and the next constraint explicit.

Governance must cross asset owners

A corridor governance charter

The corridor needs coordination across legitimate authorities

Start the governance design with the interfaces that no single asset manager controls, from eligible release to accepted delivery. Their actors retain different ownership, duties and incentives. Assign the coordination function a clear service purpose and limited powers, then preserve the distinct authority of operators, technical functions, commercial parties and public institutions. Where a proposed corridor-manager diagram implies a power, identify the agreement or law that would provide it. If that basis is absent, route the decision to the actor that holds it instead of placing the actor beneath an unsupported hierarchy.

The charter should begin with the approved service definition. It should name the participating boundaries, cargo eligibility, operating envelope and customer commitments. That gives the coordination function an object to manage. Without it, meetings can become general discussions of logistics performance in which every actor uses a different success measure. The charter should also state exclusions and unresolved external dependencies so that the coordinator's accountability matches its actual authority.

Participation should follow decisive interfaces, not prestige. A release desk can matter as much as a major equipment owner when its decision governs eligibility. A receiving customer can matter when its acceptance window constrains the service. A maintenance or infrastructure actor can matter when its calendar shapes usable capacity. The charter should identify who must attend which decision and who can provide confirmation asynchronously, avoiding a permanent committee that includes everyone but resolves nothing.

The sponsor should distinguish coordination from contracting. A governance charter can establish a process for shared planning, evidence and escalation. It cannot silently amend an access right, equipment warranty or public duty. Those changes need their relevant authority and

agreement. The charter should name that boundary rather than imply that collaboration creates consent. This protects the participants and makes the service promise more credible.

Exhibit: coordination within scoped authority

Read the relationships as declared service conditions, with the limits stated in the caption.

COORDINATION CONNECTS DECISIONS WITHOUT REPLACING AUTHORITY



Each release needs the relevant evidence and the actor authorised to decide.

Original proposed governance diagram. Connections represent coordination rather than hierarchy or an executed agreement. Technical fitness, public permission, operator deployment and commercial priority retain their relevant decision owners. The coordinator connects evidence and service consequences; it does not acquire those authorities by appearing at the centre.

Assign decisions and evidence, not only responsibilities

A responsibility matrix is useful only when it identifies an actual decision and the evidence needed to make it. The charter should state who can confirm eligible cargo, allocate a shared operating slot, suspend dispatch, approve restart and close recovery. It should distinguish recommendation, approval and execution. An actor may advise on a decision while another holds authority, and a third records the resulting event. Collapsing these roles into responsible can leave a critical handoff unresolved.

Normal planning should use a common readiness vocabulary. Forecast, committed, ready, restricted and completed states should retain their meaning across actors. The data dictionary should identify the source, update duty and correction route for each state. The coordinator should

not overwrite a technical restriction with a commercial green status. When the states disagree, the charter should define which condition governs movement and how the conflict is escalated.

An incident protocol should begin with the service impact and immediate constraints. It should identify the actor that observed the event, the technical or public stop authority, the relevant users and the current buffer or reserve state. The coordinator then aligns permitted responses. This is a different duty from determining legal liability or assigning blame. Keeping those processes separate can improve both operational candour and later evidence review.

The evidence repository should retain decisions, assumptions and revisions. A meeting minute that records only action items may lose why an alternative was rejected or which condition made approval valid. The record should include the service edition, evidence references, unresolved challenge and accountable release. It should be concise enough for operators to use but sufficiently complete for a later reviewer to trace the decision. A longer archive without decision lineage is not stronger governance.

Design escalation that can actually resolve a conflict

Escalation should follow the nature of the conflict. A missing readiness confirmation may be resolved by the interface owner. A priority conflict may need the agreed commercial decision process. A technical restriction may need qualified review. A public permission may remain with the relevant authority. The charter should route each category to an actor capable of deciding it, rather than send every issue to the most senior person and assume hierarchy supplies expertise or legal power.

The escalation record should state the feasible options and their service consequences. If a yard is full, the choices might include revised upstream release, alternate eligible storage or changed vessel allocation, each subject to its own conditions. The coordinator should preserve the effects on other users and the authority required. A decision to prioritise one commitment should be transparent within the applicable information rules, especially when it creates delay elsewhere.

Unresolved conflict should have a safe operating state. If the actors cannot agree that a movement is eligible, the charter should specify how it is held and who protects cargo and resources. Commercial urgency should not erase uncertainty about a noncompensable condition. The record should

identify the next evidence action and escalation route. This makes the hold a governed state rather than an unmanaged queue whose costs and responsibilities are later disputed.

The charter should also limit recurring exceptions. A temporary priority or recovery rule can become normal practice unless it has an expiry and review trigger. Repeated exceptions may indicate that the base service design no longer fits demand or operating conditions. The coordinator should bring that pattern into a service-basis review, not continue to resolve each episode informally. Governance then supports redesign rather than merely sustaining a fragile arrangement through individual intervention.

Review governance through service evidence

Evaluate the charter against the service decisions it is intended to resolve. Proposed evidence includes timely readiness confirmation, clear incident ownership, fewer unresolved handoff states and a traceable response to changed conditions. Define a baseline for those categories and record concurrent operational changes before interpreting any movement as governance improvement. This paper reports no measured outcomes for them. Counting meetings or participants cannot answer the service question; the review needs evidence of the relevant decision and its consequence.

The charter should include a challenge route for participants. An operator should be able to dispute a metric, capacity basis or attributed cause with evidence. A source owner should be able to correct an erroneous readiness status. A customer should be able to show that a reported delivery did not meet the agreed acceptance event. The process should record the challenge and resolution, preserving adverse findings instead of treating disagreement as obstruction.

Test the charter with a hypothetical disputed readiness message. The yard records ready space, the railway records an unreleased formation and the customer records a missed delivery. The coordinator should reconstruct which transition failed and convene the actor able to change it. If the formation lacked cargo release, the yard cannot cure the problem by accepting a queue. If the formation was eligible but the yard's readiness omitted compatible handling, the receiving definition needs correction. If both were ready but no admissible path remained, escalate the scheduling conflict. A functioning charter routes each branch to a decision; a meeting that merely repeats the competing statuses has not resolved the service interruption.

To distinguish decision latency from operating latency, record when the conflict became actionable, when the responsible actor received a complete request, when a disposition was issued and when the operating transition occurred. Delay before a complete request points to a definition or information gap; delay afterward calls for review of authority or resources. A prompt decision followed by slow execution directs the diagnosis elsewhere. These proposed qualitative diagnostics need actual observations before becoming performance findings. They give charter review a specific agenda and keep missed delivery from being attributed to bureaucracy or equipment without a traceable mechanism.

Completed proposed corridor governance charter

This is an original fictional charter structure, not an executed agreement. Each clause defines an operational purpose and identifies the approval still needed.

Completed proposed corridor governance charter			
Clause	Proposed rule	Accountable role	Required confirmation
Service purpose	Coordinate eligible producer release to accepted vessel transfer	Sponsor and service coordinator	Participating actors agree boundary and exclusions
Readiness data	Keep forecast, commitment and actual event separate	Each originating interface owner	Data definitions and permitted recipients agreed
Dispatch decision	Move only when required eligibility and downstream conditions hold	Operator within its authority	Release and access confirmations recorded
Incident coordination	Preserve stop authority and align feasible responses	Coordinator and affected owners	Technical, public and commercial powers retained
Recovery closure	Close after agreed service and reserves are restored	Incident governance roles	Evidence of stable handoffs and residual restrictions
Challenge and change	Record disputed basis and revise on decisive service changes	Model and charter owners	Challenge resolution and service edition approved

Governance operating test

Before adopting the charter, walk a normal handoff, an incident and a priority conflict through its rules. Confirm that each decision reaches an actor with actual authority and that the required evidence exists. Revise clauses that depend on assumed data access or consent. Set a trigger for service-basis review when recurring exceptions or new interfaces appear.

1. Name service boundaries and exclusions.
2. Assign forecast, approval, execution and recording separately.
3. Route conflicts by technical, commercial or public authority.
4. Retain decisions, challenges and the conditions attached to approval.

Coordination can become an expensive administrative layer

A charter should remove unresolved decisions rather than add reporting for its own sake. Limit shared data to the service purpose and convene the actors needed for the current issue. If an existing operating arrangement already resolves the relevant interfaces, document and test it instead of creating another committee. Governance quality depends on authority and usable evidence, not institutional complexity.

The charter is proposed and unexecuted

Adopt this fictional role allocation only after the relevant agreements, jurisdictional review and operational tests establish the proposed powers and data access. No actor has accepted it, and the paper creates no legal authority or data right. The charter can organise those confirmations and route an absent one to its owner; it cannot replace the acceptance needed for actual coordination.

Evidence and scope

The World Bank port-governance module describes port authorities as actors that reconcile stakeholder interests, build port communities and cooperate beyond local boundaries. It also retains the importance of legal setting and the relationship with government. This supports coordination with scoped authority; it does not establish an executed arrangement for the fictional service.

Evidence: Port Reform Toolkit: Port Governance

Decision implications

- Coordinate across boundaries while preserving legitimate authority.
- Assign actual decisions and their evidence requirements.
- Test the charter through normal, incident and conflict journeys.

GOVERNANCE AND APPLICATION

19

A commissioning plan must test the whole chain

An end-to-end commissioning worksheet

Acceptance must follow the service architecture

Commissioning should demonstrate the service that justified the investment. A component test can establish that equipment performs under specified local conditions. A system test can establish that connected components and interfaces work together. An operational service test can establish that eligible cargo reaches the chosen acceptance boundary under the declared envelope. These are complementary kinds of evidence. Treating a successful local test as full corridor acceptance leaves the dependencies that matter to the customer unexamined.

The test plan should trace each requirement to evidence and an accountable decision. A loader-rate requirement needs compatible material, measurement basis and test conditions. An end-to-end delivery requirement needs cargo release, transport, receipt, eligibility and destination acceptance events. A recovery requirement needs reserve activation and restoration. The plan should identify which requirement each test establishes, what it excludes and who can accept the result. A single commissioning percentage cannot preserve those distinctions.


The system boundary should include institutional and informational handoffs. A physical movement can succeed through manual intervention while the intended dispatch or release interface remains unready. That may be an acceptable temporary arrangement if it is safe, authorised and documented. It is not proof that the normal operating model works. The commissioning record should distinguish temporary workarounds from the commissioned service, with owners and conditions for removing or retaining them.

The relevant authorities retain their scope. A supplier can accept an equipment test; an operator can confirm readiness; a public actor can grant a permission; a customer can accept delivery. The sponsor should coordinate these decisions without pretending that one signature covers them all. An explicit acceptance matrix protects the project from both overclaiming and an unmanageable demand that a component supplier guarantee every external dependency.

Exhibit: evidence gates for the accepted service

Read the relationships as declared service conditions, with the limits stated in the caption.

ACCEPTANCE ADVANCES ONLY AS THE EVIDENCE ADVANCES



Temporary support and residual restrictions remain visible at every gate.

Supplier milestone, operational readiness and customer acceptance are distinct.

Original commissioning architecture. A component test, eligible interface, repeating resource cycle and restored reserve establish different properties. Temporary support and residual restrictions remain visible throughout. The drawing contains no completed test and no acceptance claim for an actual corridor.

Test interfaces before celebrating the first movement

The first movement is a useful milestone, but it can be specially prepared. Extra staff, prepositioned stock and exceptional priority may remove the ordinary constraints. The commissioning report should state those conditions rather than imply that the event demonstrates a sustained service. A repeatable operating test should use the declared resource and governance basis, with deviations recorded. The difference between a successful demonstration and an accepted recurring service is an important status boundary.

Interface tests should follow the actual handoff. Cargo identity and quantity should reconcile between release and receipt. Readiness messages should reach the right actor in time to support dispatch. The receiving stage should confirm that it can accept the cargo and preserve eligibility. The resource should return to its next duty when the service cycle requires it. Each test should expose the evidence gap that could prevent repeating the movement, instead of measuring only the most visible productive interval.

The programme should include adverse but admissible conditions. A changed priority, delayed upstream release or temporary restriction can reveal whether the operating plan remains coherent. Tests must remain safe and authorised; they should not manufacture a dangerous event for presentation value. A tabletop exercise can test authority and

communication, while a controlled physical exercise can test compatibility or substitution. The report should identify which property each exercise establishes and which still depends on evidence not yet available.

Failures should be classified by requirement and mechanism. A missed handoff can arise from physical capability, missing information, unconfirmed permission or unclear decision authority. The corrective action should address that mechanism and the subsequent test should show whether it was resolved. Retesting an unrelated local rate does not close an interface defect. The challenge log should retain the failed test, corrective action and evidence of closure rather than replace the record with the final favourable result.

Ramp-up needs a provisional service basis

A ramp-up plan can deliberately begin within a narrower envelope. The team may use a restricted cargo mix, temporary resource allocation or conservative dispatch sequence while it verifies interfaces. Those conditions should be labelled provisional and connected to the next acceptance gate. A lower initial rate does not necessarily show design failure, just as a high isolated rate does not prove full readiness. The review should compare the evidence with the staged service requirement.

The plan should state how it advances. Each release should require relevant evidence such as stable receipt, reconciled quantity, valid permissions or restored reserves. An elapsed date alone is a weak readiness criterion when the unresolved dependency remains. Conversely, excessive ceremony can delay service after the necessary evidence is established. The gate should be specific enough that operators and reviewers can recognise when the condition has been met.

Temporary support needs a sunset or adoption decision. Additional commissioning personnel, manual interfaces or dedicated equipment can become permanent without their costs and authority being reviewed. The service basis should identify what normal operation assumes and what ramp-up support supplies. If the project chooses to retain a workaround, revise the operating and commercial basis explicitly. A silent continuation can leave the accepted annual capacity dependent on resources absent from the approved plan.

Sustained operation should also test restoration of readiness. A chain that delivers a staged campaign while depleting buffers or consuming exceptional fleet reserve may not be able to repeat it. The commissioning record should show whether the intended stock, equipment and

operating positions are restored. This links the acceptance test to the maintenance and buffer chapters: delivery, recovery and readiness for the next cycle are separate properties that the service promise may require.

A repeated cycle challenges a staged demonstration

A fictional first delivery is prepared with fully released cargo, a dedicated path, extra staff and empty receiving space. It succeeds and provides valuable evidence of compatibility. The next service has to share resources, form cargo through normal release and return equipment through the ordinary terminal process. It may reveal a circulation or readiness problem absent from the demonstration. The commissioning record should therefore state the exceptional support and identify the repeating-cycle test that remains.

The test should trace the resource back to its next eligible assignment. A delivery that immobilises equipment at destination can succeed at the cargo boundary and fail the fleet-circulation requirement. A campaign that consumes the protective yard reserve can succeed for the current vessel and leave the next commitment exposed. These are distinct outcomes. The acceptance matrix should preserve both rather than force a single success label across the whole programme.

If the repeat test fails, the corrective action should address the limiting state. A delayed return may require terminal release or resource allocation changes rather than a higher loading rate. An exhausted reserve may require a restoration sequence rather than more isolated peak capability. The evidence should identify whether the original annual assumptions remain defensible and whether the approved intervention addressed the right mechanism. The project should not close a service defect through an unrelated component test.

A provisional operating state can be appropriate while these issues are resolved. The sponsor should state the permitted duty, temporary support, unresolved requirement and evidence needed to advance. That allows useful service to begin within a controlled envelope without claiming a full capability that has not been demonstrated. It also protects the local equipment acceptance from being confused with wider readiness, giving each actor a clear and feasible evidential responsibility.

Close acceptance with residual conditions visible

A residual defect can be minor for one duty and blocking for another. The acceptance decision should state which service is permitted, under which restrictions and with which open actions. It should not use a generic accepted status that obscures a consequential limitation. An unresolved data defect may prevent a performance claim even when safe movement is possible. An unresolved technical restriction may prevent movement regardless of a favourable data record. The matrix should preserve the difference.

The committee should receive evidence of the delivered system, not only progress against the commissioning schedule. An integrated schedule helps identify dependencies and prepare tests, while completion of an activity does not prove that its requirement was satisfied. The acceptance review should connect the schedule, test record, defects and actual authority to release. This prevents a programme marked complete from becoming an unsupported service certificate.

A commissioning result can update only the capacity property it actually tests. A hypothetical repeated-cycle demonstration that relies on dedicated paths supports circulation under that path arrangement; it does not establish capacity under the ordinary shared timetable. A successful loading test with prepositioned eligible stock supports the loader duty but leaves replenishment untested. Preserve those test conditions in the basis bridge and mark the unresolved transition that prevents their combination into an annual guarantee. Where ordinary operation reproduces the conditions, the bridge can advance. Where it does not, the next test should challenge that difference rather than repeat the favourable demonstration.

The residual disposition should distinguish a restriction from a performance-evidence gap. A hypothetical cargo trace can be incomplete while movement remains technically permissible; acceptance of a throughput claim must then wait for lineage reconciliation. A handling restriction can remain even when every timestamp is complete; movement must stay within the permitted envelope. A temporary dedicated crew can permit provisional service while leaving the normal staffing duty unresolved. Recording these different residuals tells the operator which service can continue and tells the sponsor which claim remains unsupported. It avoids both stopping safe permitted work merely because a report is unfinished and declaring full service acceptance merely because the paperwork is complete.

Completed illustrative commissioning matrix

This proposed record contains no executed test or acceptance claim. It links each service requirement to evidence and a release decision.

Completed illustrative commissioning matrix			
Requirement	Test or review	Evidence still needed	Acceptance role
Local loading capability	Specified cargo and equipment test	Rate record, conditions and technical result	Supplier and relevant technical actor
Eligible cargo handoff	Trace release through yard receipt and loading assignment	Identity, quantity and status reconciliation	Producing and receiving operators
Recurring transport cycle	Repeat compatible dispatch and resource return	Path, terminal and circulation records	Transport and terminal actors
Restriction response	Safe authorised exercise of stop and restart process	State messages, authority and permitted actions	Relevant stop authority and coordinator
Reserve restoration	Verify depletion, deployment and recovery sequence	Eligible reserve and repeated-service evidence	Service and reserve owners
Whole-chain service	Accepted delivery under declared operating basis	Customer boundary record and residual defects	Sponsor, operator and receiving acceptance roles

Service acceptance gate

Accept only the service state demonstrated by relevant tests and authorised reviews. Keep component, interface and recurring-service evidence separate. Record temporary support and residual restrictions. Revise the model basis when observed operation contradicts its assumptions, and close each defect through evidence tied to the original requirement.

- 1. Trace requirements to tests and acceptance authority.
- 2. Distinguish demonstration from normal repeating operation.
- 3. Test safe adverse states and reserve restoration.
- 4. Publish the permitted envelope and open conditions at handover.

Full recurring-service proof may delay contractual acceptance

Contractual component acceptance and sponsor service acceptance can be separate decisions with different evidence. The paper does not recommend withholding a supplier's agreed milestone because of an unrelated external dependency. It recommends preserving what that milestone establishes and testing the corridor service through the appropriate programme. The sponsor should allocate those duties in advance so that a technical test does not become an improvised dispute over wider responsibility.

The matrix records a proposed test plan

Before accepting performance, legal readiness or customer delivery, obtain the executed test and scoped acceptance evidence from the relevant actors. No commissioning event has been executed for this fictional corridor. The matrix supplies a structure for documenting those properties when suitable evidence exists; its entries alone cannot certify them. Retain any property whose evidence remains absent as an open condition in the permitted service state.

Evidence and scope

NASA's systems engineering guidance connects requirements, technical management and decision rationale. Its reference architecture requires adaptation; it does not accept this corridor.

Evidence: NASA Systems Engineering Handbook

Evidence and scope

GAO's schedule guide treats schedules as integrated dependency models. Schedule completion alone is not evidence that an operational service requirement has passed its test.

Evidence: Schedule Assessment Guide: Best Practices for Project Schedules

Decision implications

- Accept the service demonstrated by the evidence.
- Keep local equipment tests distinct from whole-chain delivery tests.
- Preserve provisional states and residual restrictions at handover.

When the bottleneck thesis is incomplete

An adversarial model-validity map

The thesis is conditional on the objective

The bottleneck thesis is useful when the objective is increased accepted serial flow under a compatible annual basis. It becomes incomplete when the decision concerns another duty or when the serial representation excludes a decisive mechanism. This chapter makes those challenges explicit. A strong thesis should identify the conditions under which it ceases to select the action, rather than protect itself by calling every adverse outcome another bottleneck after the fact.

Demand is the clearest structural challenge. If eligible demand is below every stage capacity, additional capability cannot increase the common flow in the current model. The changed-demand study demonstrates that reversal. The correct next question may concern customers, service quality, access or cost rather than another asset. Calling demand a bottleneck can be descriptively convenient, but it should not obscure that the relevant intervention is commercial and outside the app's capacity-only expansion tests.

Quality can change the objective from gross movement to accepted output. A higher rate may increase damage, contamination or rework. A stage that governs accepted quality can matter even when its gross annual capacity is high. The service map should therefore define eligibility and acceptance before the capacity comparison. The current app assumes uniform quantity and no loss, so it cannot calculate that quality tradeoff. A valid challenge requires a method and evidence representing the actual transformation or rejection mechanism.

Strategic reserve can justify capacity that does not increase normal annual flow. A sponsor may value continuity during an interruption, optionality for a future service or compliance with an access duty. Those reasons can be legitimate, but they need a stated objective and feasibility evidence. The annual minimum result does not evaluate them. The committee should avoid forcing every benefit into throughput and should avoid treating unquantified strategic language as automatic justification for any expansion.

Topology changes the governing mathematics

A serial chain requires every tonne to pass through every represented stage. A parallel network allows route choices and may contain shared links. The constraint can then be a set of links or an allocation rather than the smallest individual stage. Treating alternatives as serial stages can make more options appear to reduce capacity, while summing all alternatives can overstate capacity when they share a decisive link. Both errors arise from a mismatch between topology and method.

When routes become parallel, specify origins, destinations, capacities and shared dependencies together with access conditions and assignment rules. Physical route options become service only through that allocation basis. The present app does not implement a network model, so renaming its serial stages cannot establish network optimisation. Select and implement a suitable method whose inputs, outputs and limits can be reproduced and reviewed, or retain the current calculation as a narrower serial screening result. The representation must change before its conclusions are widened to the new network decision.

Multi-product operations can also create allocation problems. Different cargo classes may consume handling time, space and access differently. A single tonne basis can conceal the tradeoff between them. A product mix change may alter effective capacity without any physical asset changing. The team should determine whether a homogeneous annual flow remains useful for the decision or whether explicit resource consumption and allocation are needed. The study offers no calculated mixed-product optimum.

Shared resources can make apparently serial stages interact. Receiving and reclaim may use the same handling equipment; loaded and empty movements may share infrastructure; documentary reviews may use a common service desk. Assigning separate independent capacities can overstate what can be achieved simultaneously. The model needs to represent the shared constraint or narrow its claim to a screening bound. The challenge is methodological, not a request to insert another arbitrary availability reduction until the result looks plausible.

Time and noncompensable conditions can dominate

The annual minimum can be unattainable when service opportunities do not align and buffers are finite. A stage can have ample annual capability but be unavailable during the only compatible downstream window. A yard can fill and block the upstream chain despite adequate annual receipt

and loading rates. The remedy may involve chronology, batching or buffer deployment rather than another annual-capacity increase. A suitable time-based study should represent those mechanisms when they can reverse the decision.

Establish route authority, configuration safety and cargo acceptance before ranking flow alternatives. These noncompensable conditions define the feasible set: greater throughput cannot offset an unauthorised route, unsafe configuration or unacceptable cargo. The relevant actors must confirm eligibility and authority. The current model assumes an eligible fictional service and supplies no legal or technical admissibility decision. If a condition is unresolved, keep that option outside a commitment ranking until the required review resolves it rather than convert the condition into a compensable throughput penalty.

Cost and implementation can make the physical optimum unattractive. A large coordinated expansion may create more annual capacity but require expensive work, disruptive transition or unsupported customer demand. A smaller reversible change may be preferable under the actual economic basis. The canonical cases deliberately isolate physical response and operating denominator effects. They do not establish the best capital programme. The paper's recommendation remains to use the system constraint as part of the decision, not as its sole criterion.

Evidence can also be the binding decision issue. If the team does not know whether observed rail shortfall reflects capability, demand or destination blocking, a precise expansion calculation cannot resolve the uncertainty. A focused measurement may be more valuable than another scenario. The analysis should identify the observation that would distinguish the competing mechanisms and the consequence for the choice. This is a limit on what the current evidence permits, not a claim that every decision must wait for perfect information.

An objective change can preserve the asset and reverse the recommendation

Suppose the fictional loading-only intervention is proposed to increase annual corridor quantity. The frozen result does not support that claim because rail remains binding. Now change the stated objective to completing a ready vessel transfer within a shorter permitted operating window. The same physical asset proposal may become relevant, but the annual app cannot establish its value. The correct response is a new service question and suitable evidence, not a defence that the original throughput model secretly represented the new objective.

This distinction permits serious disagreement without forcing a false binary. The annual result can be correct and the proposal can still deserve further study. The proposal can also be attractive under a local service objective while failing a wider cost or access condition. The committee should identify which proposition it is deciding and which evidence supports it. Changing the objective is legitimate when declared; changing it after a failed test while preserving the original success claim is not a transparent evaluation.

The reverse case also matters. A coordinated expansion can raise the fictional annual bound while worsening a service objective omitted from the model. It could create a difficult transition, increase exposed inventory or require an unacceptable access arrangement. No such consequence is asserted as an actual finding here. The point is that a positive physical output does not automatically dominate every feasible alternative. The adversarial review should test the omitted consequence most capable of reversing the choice.

The resulting review record should state whether the challenge changes the objective, feasible set, input basis or method. That classification tells the team what work is required. A corrected objective needs a revised service and evaluation plan. A changed feasible set needs relevant authority evidence. A changed input needs reproduction. A changed mechanism needs another implemented method or a narrower claim. This is more useful than adding a general caveat and leaving the recommendation unchanged.

Turn the adversary into a review record

A reviewer should challenge the thesis through an admissible counterexample. State the service, changed condition and expected effect, then explain why the annual bottleneck recommendation would be inferior or incomplete. A vague objection that reality is complicated is less useful than a specific case of demand limitation, tied constraints, shared topology, quality loss or incompatible windows. The review should distinguish a failure of the formula from a failure of its applicability to the decision.

The frozen slate supports several of these structural challenges. Demand limitation removes the capacity-expansion rationale for increased flow. Equal limits make isolated relief inadequate. An unavailable essential stage stops serial service and makes unit cost undefined. The coordinated case shows that the next constraint can govern the benefit. The cases do not cover network routing, stochastic queues or quality. Their absence is disclosed rather than filled with an unimplemented calculation.

The response should be proportionate. If the challenge identifies a basis error, correct and reproduce the model. If it identifies an omitted mechanism, select a method that represents it or narrow the recommendation. If it identifies an ineligible option, remove it from the feasible comparison. If it identifies an economic omission, obtain the missing appraisal evidence. A defence that simply adds caveats without changing the decision logic is insufficient when the challenge can reverse the choice.

A validity map should preserve the useful calculation after rejecting its original recommendation. If a hypothetical loader proposal is justified by berth-window recovery rather than annual quantity, retain the unchanged serial-flow result and move the window question to a dated operating study. If a parallel route is introduced, retain the serial calculation only for the route it represents and model the shared interfaces separately. If cargo rejection matters, retain gross movement as a handling measure and reconstruct accepted service. Each branch changes the object of inference rather than adding another disclaimer. The committee can then see exactly which conclusion survives, which is withdrawn and which new evidence would support the revised decision.

Exhibit: conditions that change the decision method

Original adversarial map. It identifies required responses and does not calculate the omitted methods.

Exhibit: conditions that change the decision method			
Challenge	Why the annual rule is incomplete	Evidence or method needed	Decision response
Demand below capacity	Physical expansion cannot create eligible demand	Demand and service-value evidence	Reframe the commercial question
Parallel routes or shared links	Single serial minimum misstates topology	Implemented network and access model	Represent the feasible paths and common constraints
Quality loss or rework	Gross tonnes differ from accepted output	Quality and transformation evidence	Change the quantity and acceptance basis
Finite buffers and windows	Annual opportunities may not align	Implemented time-based operating analysis	Test chronology and usable reserve
Noncompensable duty	More flow cannot offset ineligibility	Relevant technical and jurisdictional review	Gate feasibility before preference
Economic or implementation burden	Physical ranking omits full consequences	Incremental cost, timing and user impacts	Compare feasible alternatives on the full basis

Adversarial validity review

Require the reviewer to name an admissible condition that can change the choice. Distinguish computation error, basis error and omitted mechanism. Revise the model, method or scope accordingly. Preserve challenges and resolutions in the edition record, and leave an unresolved consequential objection open rather than claim completion from mechanical reproduction.

1. State the counterexample and service boundary.
2. Identify which assumption or objective changes.
3. Test the effect using an implemented suitable method or narrow the claim.
4. Record the decision consequence and remaining review gap.

The conditional thesis may appear too modest

Its modesty is deliberate and useful. An unqualified statement that the smallest asset should always be expanded would fail under demand, topology and service-quality challenges. The paper instead offers a reproducible annual mechanism and a practical route for identifying when another method is needed. That conditional result can guide real work without claiming that a screening formula replaces operational, commercial or jurisdictional judgement.

Mechanical depth is not substantive acceptance

Submit the argument's originality, operational soundness and usefulness to the intended audiences for independent domain review, and submit the actual pages for independent visual review. Word counts, model reproduction and page geometry establish their checked artifact properties only. Use those receipts alongside the reviews to identify a repair; do not infer acceptance of the reasoning from them or assign an author acceptance grade.

Evidence and scope

The Magenta Book links evaluation to decision needs and credible evidence. A monitored change alone does not establish a causal intervention effect.

Evidence: Magenta Book: Central Government guidance on evaluation

Decision implications

- Test the thesis against conditions that can reverse the choice.
- Change the method when topology, chronology or quality demands it.
- Keep independent substantive review separate from mechanical checks.

PRACTICAL APPENDICES

21

Appendix: capacity and flow dictionary

A canonical app replication dictionary

The dictionary is part of the reproducible model

The replication dictionary defines what enters and leaves the annual Corridor lab computation. Its purpose is to let a reader reconstruct the fictional studies and identify where an actual application would need more evidence. It is not a glossary of every logistics term. Each entry connects a quantity or state to the method, its unit and its decision role. A reader should be able to explain why a value belongs in the annual minimum rather than merely recognise its label.

Demand is the chosen annual quantity of eligible service in the fictional chain. It is not a sales forecast, production aspiration or measured market estimate. The app uses it as an upper bound on common flow. A lower demand fixture can therefore make the chain demand-limited without changing any asset. For an actual application, the analyst should define what qualifies as eligible demand, whether it is contracted or assumed and which acceptance boundary the quantity refers to.

A stage is a required serial function through which all modelled cargo passes. Its name identifies that function within the fixture. Adding an alternative route as another stage changes the meaning incorrectly because the app will require the common flow to pass through it as well. A stage should not be inserted merely to represent an interesting stakeholder. Its capacity input must correspond to a required function that can constrain the chosen serial service.

Nominal capacity is an annual rate before the fixture's explicit availability multiplier. It uses the same cargo quantity basis across stages. It should not be confused with stored material, fleet count or instantaneous equipment rate. Any derivation from those objects lies outside the current app and needs its own validated method. The dictionary therefore makes the annual convention visible instead of pretending that a unit label proves how the input was obtained.

Availability and costs have narrow conventions

Availability is a percentage input that the engine applies to nominal capacity to produce effective annual capacity. It represents the fictional annual usable-capacity assumption. It does not establish a probability distribution, mechanical reliability or event timing. The input definition should avoid double-counting downtime already included in the nominal annual rate. For actual use, the team should state whether the percentage includes maintenance, crew, access or other readiness conditions and preserve that basis with the scenario.

Variable cost is the assumed operating amount per common tonne attributed to a stage. The engine sums those stage inputs. It does not calculate a supplier tariff, marginal route response or changes in cost at different operating rates. Fixed cost is the assumed annual operating obligation attributed to a stage. The engine sums it and allocates it to common throughput when that throughput is positive. The fictional costs are illustrative and represent no client, transaction or market quotation.

The quantity and currency conventions make that allocation dimensionally coherent. Annual fixed amounts and annual throughput use matching million-scale units, yielding the labelled cost per tonne. The manuscript prints canonical scalar outputs rather than recomputing them in authoring code. This prevents a separate calculation from drifting away from the public app. It does not expand the model's economic scope: capital timing, financing, taxes, replacement and residual value are absent unless separately represented by a suitable method.

Undefined unit cost is an important output state. When common throughput is zero, the engine cannot divide fixed obligations by a positive delivered quantity, so it reports an unavailable unit cost. Fixed obligations remain in the fixture. This should not be displayed as free service, a finite average or a numerical placeholder that invites comparison with positive-flow cases. The reader should instead inspect the zero-service condition and the relevant continuing obligations.

Outputs describe a common annual balance

Throughput is the minimum of eligible demand and all effective stage capacities. It is a calculated annual quantity within the serial convention. The output is not an observed delivery total. It does not establish how work is scheduled, whether buffers are sufficient or whether cargo is

accepted at every actual interface. Those are applicability conditions that the main chapters require the practitioner to review before treating a real input slate as a service promise.

A bottleneck output lists stage names whose effective capacity equals the common flow under the engine's comparison rule. Demand limitation can produce an empty list when every stage has capacity above demand. Equal capacities can produce several names. The interpretation should preserve both states. An empty list does not mean the service has no problem; it means the represented annual capacity stages do not bind under the chosen demand. Several names imply that isolated relief may not create common-flow gain.

Stage utilisation compares common annual flow with that stage's effective annual capacity. It is unavailable for an effective capacity that is zero. Spare capacity is the difference between effective annual capacity and common flow. These are model-balance quantities, not measured shift utilisation, stock reserves or guaranteed recovery margins. Their useful role is to show which represented capacities exceed the current annual bound and to prompt review of why those margins exist.

Stage annual cost combines its fixed operating input with variable cost at common flow. Stage unit cost allocates those amounts to common positive throughput. Whole-chain unit cost combines the stage costs using the same flow. The model assumes every tonne incurs every serial stage's variable cost. A service with losses, bypasses, transformations or rework needs another quantity and cost representation. The dictionary helps a reader identify that mismatch before accepting the displayed totals.

Expansion and stress are tests with declared constants

The app's single-stage expansion tests change one nominal capacity according to its implemented method and hold availability, demand, other stages and costs fixed. Each result reports the new effective capacity, common throughput, gain and next constraint. These are deterministic sensitivity tests. They are not procurement designs and contain no investment cost. The manuscript's frozen cases also include explicitly authored interventions, whose complete inputs are retained separately from these built-in tests.

The availability stress changes each stage's availability according to the implemented app rule, with a lower bound preventing a negative capacity assumption. It then calculates stressed common flow and operating unit cost. The result is a scenario, not a likelihood-weighted forecast. Its value

is to expose dependence on the annual readiness inputs. A practitioner should not infer a confidence level or treat the fixed stress as covering every relevant interruption mechanism.

The fixture files preserve complete inputs, and the output files preserve the canonical result object, including tables, methods, warnings and limitations. Bindings identify scalar input or output paths for every printed model quantity. Claims retain the illustrative basis, source fixture, sensitivity and manuscript block references. This chain allows the audit to execute the current engine, compare exact outputs and resolve each published token without a parallel numerical implementation.

The dictionary can diagnose a reuse error through the identity being transferred. A nominal annual rate copied from a supplier test needs an availability and calendar bridge. A storage volume copied into a stage-capacity field needs a turnover mechanism; changing its unit label cannot supply one. A route alternative represented as another mandatory stage changes parallel choice into serial obligation. An undefined unit cost copied as a neutral amount erases the zero-flow denominator. These are different repairs, and none requires disputing the engine's arithmetic. Applying this translation test before computation prevents a faithful calculation from silently answering a different physical question than the one the analyst intended.

Exhibit: canonical replication dictionary

Original input and output map. All numerical examples elsewhere use frozen scalar bindings; this table states concepts and units only.

Exhibit: canonical replication dictionary			
Field	Unit or state	Canonical role	Interpretation limit
demand	Mt/year	Upper bound on common annual flow	Chosen scenario, not market evidence
stages.capacity	Mt/year	Nominal annual stage rate	Needs an external derivation for real use
stages.availability	Percent	Multiplier for effective capacity	No failure chronology or probability
stages.variableCost	USD/t	Operating cost on every common tonne	No tariff or marginal-response model
stages.fixedCost	USDm/year	Annual obligation allocated to flow	Expansion capital absent
throughput	Mt/year	Minimum annual bound	No dispatch or inventory simulation
bottlenecks	Stage names or empty list	Represented binding capacity stages	Demand can bind without a named stage
unitCost	USD/t or unavailable	Whole-chain operating allocation	Undefined at zero flow; no investment return

Scenario reuse gate

Before reusing a fixture, verify that the service remains serial, the cargo basis is uniform and the annual inputs have the same meaning. Preserve method version, complete inputs and output bindings. Reproduce after a consequential change and review whether the changed method or service requires another model. Keep fictional scenarios explicitly illustrative in every exported use.

1. Read units and limits before editing an input.
2. Trace printed quantities to scalar bindings and complete fixtures.
3. Reject routing, timing or quality claims the engine does not calculate.
4. Record revisions and rerun exact reproduction.

A detailed dictionary can discourage tool adoption

The user does not need to read the entire appendix before trying an illustration. The interface can expose concise methods and limits alongside useful defaults. The detailed dictionary becomes essential when a result is exported, reviewed or reused for a consequential decision. Its purpose is to support trustworthy reuse, while the app remains usable as a transparent learning instrument.

A reproducible fixture is not verified operational evidence

The engine can execute the frozen inputs correctly and the ledger can bind every printed scalar. That establishes reproducibility. It does not establish that the annual inputs were measured, that the operation is feasible or that the service will be delivered. Those properties need independent relevant evidence.

Decision implications

- Read each field as a convention with a decision role.
- Preserve undefined and demand-limited output states.
- Reproduce changed fixtures and reject claims outside the implemented method.

PRACTICAL APPENDICES

22

Appendix: illustrative corridor studies

Reproducible scenarios and full outputs

Read the studies as controlled comparisons

The study slate is expressly fictional. It uses named serial functions to make the mechanism easy to follow, not to imply an actual mine, railway or port. Complete inputs were authored before freezing and executed through the Corridor lab compute function. The output objects, scalar bindings and claim ledger are retained with the paper. The manuscript does not duplicate the numerical method in its authoring code. This appendix describes what changes, what remains constant and how each case challenges the service decision.

The common base has producer release, rail movement, port stockyard and vessel loading. All stages use the same illustrative cargo quantity and incur their stated variable and fixed operating costs on the common flow. Demand is a chosen annual bound. The stage capacities and availability assumptions are chosen to create a rail constraint close to the loading

limit. That structure makes the next-constraint effect visible without assigning a measured frequency or claiming that a particular industrial geography behaves this way.

Every intervention case preserves the declared serial topology and homogeneous cargo basis. Operating costs remain unchanged across the physical interventions to isolate capacity and denominator effects. This is a deliberate analytical restriction. It does not imply costless expansion. A capital appraisal would need the incremental investment, implementation period, additional operating obligations and customer value. The slate should therefore be used to understand the physical method and its applicability limits, not to select an actual supplier or investment programme.

The additional cases challenge the original interpretation rather than simply extend a preferred expansion story. Lower availability tests readiness sensitivity; lower demand removes the annual capacity bottleneck; tied capacities challenge isolated relief; zero loading availability tests the undefined-cost boundary. Their purpose is to give an independent reader admissible counterexamples within the implemented method. Other important challenges, such as routing, queues and quality, remain outside the engine and are handled conceptually in the main analysis.

The base, loading and rail cases isolate the next limit

Illustrative calculation. Assumptions and sensitivity are published with the model.

The base chain calculates 11.90 Mt/year and 23.73 USD/t, with Rail corridor binding. The capacities, availability and cost inputs are printed below on their declared units. The result is a common annual flow, not the sum of local outputs. Its main implication is that every stage carries the same modelled quantity while the smallest effective annual bound prevents the higher-capacity stages from increasing that quantity on their own.

Illustrative. The loading-only case changes the vessel-loading nominal capacity and nothing else. It calculates 11.90 Mt/year and 23.73 USD/t, with Rail corridor still binding. The unchanged common result falsifies an unqualified claim that enlarging that visible asset must raise annual corridor flow. It does not falsify a separate marine-window or resilience benefit, because those outcomes are absent from the method. The reader should preserve that narrow scope when using the comparison.

Illustrative. The rail-only case changes the original binding nominal capacity and calculates 12.00 Mt/year and 23.67 USD/t. The next annual constraint is Vessel loading. The implication is not that rail relief is pointless. It is that the represented gain is limited by another unchanged stage. A design brief should therefore include that stage's readiness and any relevant service mechanism before the sponsor interprets the nominal rail increment as added delivered quantity.

The built-in single-stage expansion table offers a related controlled test for the base fixture. Its rows change one stage according to the implemented method, preserving all other inputs. The manuscript publishes the resulting scalar readings and keeps the complete table in the frozen output. This is useful for checking the logic of isolated relief, but it still omits feasibility and investment costs. A positive tested gain identifies a physical opportunity under the model, while a zero gain identifies an annual limitation rather than a universal absence of value.

Coordinated relief and lower availability change different mechanisms

Illustrative calculation. Assumptions and sensitivity are published with the model.

The coordinated case changes rail and vessel-loading nominal capacities together. It calculates 14.40 Mt/year and 22.39 USD/t, with Port stockyard binding. The result demonstrates that joint relief can reach a different annual bound from either isolated intervention. It does not establish that simultaneous procurement is superior to a staged operating and investment sequence. That choice depends on evidence outside the fictional operating-cost comparison.

Illustrative. The lower-availability case preserves nominal capacities and costs but changes the rail readiness assumption. It calculates 7.00 Mt/year and 29.14 USD/t, with Rail corridor binding. This case isolates the distinction between physical nameplate and assumed usable annual capacity. Its sensitivity is structural: less usable capability at the limiting stage reduces common flow and raises the fixed-cost allocation per delivered tonne under the unchanged cost basis.

Illustrative. The engine's deterministic stress produces a further result for each case. The base stressed flow is 10.50 Mt/year; the coordinated stressed flow is 12.80 Mt/year; the lower-availability stressed flow is 5.60 Mt/year. These quantities are scenario outputs, with the stress rule retained in the canonical methods metadata. They do not have likelihoods and should not be presented as expected outcomes or confidence bounds.

The comparison also clarifies which evidence a practitioner would need. Coordinated physical relief needs technical and operating feasibility at several stages. An availability improvement needs a defined readiness mechanism and proof that the changed calendar or capability is usable. Both need demand and acceptance evidence before becoming a service promise. The frozen cases cannot supply those properties, but they prevent the committee from confusing a nominal expansion, a readiness change and a demonstrated increase in accepted service.

Demand, ties and zero flow supply boundary challenges

Illustrative calculation. Assumptions and sensitivity are published with the model.

The lower-demand case changes only the annual demand bound. It calculates 8.00 Mt/year and 27.50 USD/t, with no capacity stage listed as binding. The result reverses the capacity-expansion interpretation because the customer-side bound is below every effective stage capacity. A practitioner should then examine eligible demand and the service proposition rather than claim that physical additions create more current flow. Spare capability can still serve a separate reserve objective if its duty is defined.

Illustrative. The tied case changes rail's nominal and availability basis so that its effective annual capacity matches loading. It calculates 12.00 Mt/year with Rail corridor and Vessel loading binding. The annual model warns that isolated relief may produce no common-flow gain. This is an admissible counterexample to a simplistic single-asset recommendation. A real application should investigate whether equal limits are genuine or merely an artefact of rounded or poorly defined capacity inputs.

Illustrative. The zero-service case sets loading availability to 0.00 percent and calculates 0.00 Mt/year. Whole-chain unit cost is not defined. The essential serial function is unavailable, so upstream spare capacity cannot substitute for it within this topology. Fixed operating obligations remain present. The output should be interpreted as an unavailable per-tonne ratio, not as a finite estimate of the financial loss or as evidence that the continuing obligations vanish.

These cases should be retained in review even if an investment narrative prefers the expansion cases. They demonstrate why the method needs a demand definition, correct handling of tied constraints and a clear unavailable state. They also show the limits of the slate: an alternate loading route would change topology and require another method. The

zero-service result cannot establish that no real workaround exists; it establishes that the declared essential function has no substitute in this particular serial representation.

Reproduce the complete object and challenge its scope

A replication should begin with the exact frozen input, not with numbers copied from a rounded PDF table. The table supports reading and audit, while the JSON preserves full input and output precision. Execute the canonical compute function and compare the complete stable result object with the frozen output. A difference can reflect a fixture change, an engine change or an error. The reviewer should identify which occurred before interpreting the revised conclusion.

Scalar bindings use explicit input or output paths. A printed token resolves to one scalar in the canonical object and to one ledger claim. The claim records illustrative status, method, fixture and sensitivity, with reciprocal manuscript references. String readings such as binding-stage names are included, so directional conclusions are tied to the same reproduced object as the numerical values. This prevents a narrative from retaining an outdated bottleneck label after a quantity or method changes.

The complete output also retains tables, warnings, methods and limitations. Those fields are part of the interpretive record, even when the PDF prints only the decision-relevant scalars. A reviewer should inspect them alongside the stage results and intervention tests. Warnings about demand, shared constraints or zero flow should not be discarded because they make the headline less simple. A faithful export should preserve the method and limits as well as the attractive result.

The study slate also supports a contrast-selection instrument. Use loading-only against base to test an intervention at a nonbinding stage; use rail-only against base to test relief of the original limit; use coordinated against rail-only to expose the next receiving duty. Use lower availability to separate readiness from nominal rating, lower demand to separate use from capability, tied limits to test a joint constraint and zero service to test an undefined denominator. Selecting the pair by the disputed mechanism gives replication a practical purpose. Combining all changed cases into an apparent trend would lose that purpose because the fixtures alter different causes. These remain fictional controlled contrasts, with no field validation supplied.

Illustrative calculation. Assumptions and sensitivity are published with the model.

Exhibit: physical inputs for the original study cohort

Illustrative annual inputs. Stage cells state nominal capacity in million tonnes per year followed by availability in percent. Names and operating cost inputs remain common across the slate. Additional numerical boundary studies are reported separately.

Exhibit: physical inputs for the original study cohort					
Study	Demand, Mt/year	Mine rate / availability	Rail rate / availability	Yard rate / availability	Loading rate / availability
Base	20.00	18.00 / 90.00	14.00 / 85.00	16.00 / 90.00	15.00 / 80.00
Loading only	20.00	18.00 / 90.00	14.00 / 85.00	16.00 / 90.00	20.00 / 80.00
Rail only	20.00	18.00 / 90.00	20.00 / 85.00	16.00 / 90.00	15.00 / 80.00
Coordinated	20.00	18.00 / 90.00	20.00 / 85.00	16.00 / 90.00	20.00 / 80.00
Low availability	20.00	18.00 / 90.00	14.00 / 50.00	16.00 / 90.00	15.00 / 80.00
Lower demand	8.00	18.00 / 90.00	14.00 / 85.00	16.00 / 90.00	15.00 / 80.00
Tied limits	20.00	18.00 / 90.00	15.00 / 80.00	16.00 / 90.00	15.00 / 80.00
Zero service	20.00	18.00 / 90.00	14.00 / 85.00	16.00 / 90.00	15.00 / 0.00

Illustrative calculation. Assumptions and sensitivity are published with the model.

Exhibit: common illustrative operating cost inputs

The following input costs are unchanged in every study. They exclude intervention capital and any added operating obligation; the comparison isolates the throughput denominator.

Exhibit: common illustrative operating cost inputs			
Stage	Variable cost, USD/t	Annual fixed cost, USDm	Status
Mine release	2.40	14.00	Chosen illustration
Rail corridor	7.60	42.00	Chosen illustration
Port stockyard	2.80	20.00	Chosen illustration
Vessel loading	3.20	16.00	Chosen illustration

Illustrative calculation. Assumptions and sensitivity are published with the model.

Exhibit: decision-relevant canonical outputs

All displayed results are canonical scalar bindings. Complete stage balances, expansion rows, warnings and metadata remain in each frozen output JSON.

Exhibit: decision-relevant canonical outputs					
Study	Flow, Mt/year	Cost, USD/t	Availability stress flow, Mt/year	Stress cost, USD/t	
Base	11.90	23.73	10.50	24.76	
Loading only	11.90	23.73	10.50	24.76	
Rail only	12.00	23.67	10.50	24.76	
Coordinated	14.40	22.39	12.80	23.19	
Low availability	7.00	29.14	5.60	32.43	
Lower demand	8.00	27.50	8.00	27.50	
Tied limits	12.00	23.67	10.50	24.76	
Zero service	0.00	not defined	0.00	not defined	

Illustrative calculation. Assumptions and sensitivity are published with the model.

Exhibit: built-in isolated expansion tests for the base

The implemented expansion rule is retained verbatim in the canonical method metadata: Each expansion case increases one nominal capacity by 20%; availability stress reduces every stage by 10 percentage points. Each row below preserves other inputs and reports only the tested physical response.

Exhibit: built-in isolated expansion tests for the base					
Changed stage	New effective rate, Mt/year	Flow, Mt/year	Gain, Mt/year	Next constraint	
Mine release	19.44	11.90	0.00	Rail corridor	
Rail corridor	14.28	12.00	0.10	Vessel loading	
Port stockyard	17.28	11.90	0.00	Rail corridor	
Vessel loading	14.40	11.90	0.00	Rail corridor	

Frozen-study review procedure

Open the complete input and output objects for the named study, execute the canonical engine and compare exact results. Review the scalar bindings, narrative reading and warnings together. Challenge the unchanged-cost convention before drawing any economic conclusion. Record an engine or fixture change as a new reproduction event with its decision consequence.

The studies remain fictional in every reuse

The fixture slate represents no client, location or measured operation. Its cases can demonstrate a method and structural reversals, but cannot establish actual capacity, event likelihood, legal readiness or investment value. Any application needs a new evidence basis and relevant review.

Numerical equality and a supported change of scale

Illustrative calculation. Assumptions and sensitivity are published with the model.

Binding membership is a property of the supplied capacity ordering, not a convenience tolerance measured in one chosen unit. The small-distinct-capacities study gives stage A nominal capacity 1e-09 and stage B 2e-09, each at 10% availability, against demand 1e-09. Its effective capacities are 1e-10 and 2e-10 Mt/year. Only A binds. Its tested single-stage expansion leaves A as the next constraint. This artificial scale has no operating significance; it makes an incorrect absolute comparison unmistakable.

The revised engine uses rational products and comparisons for the canonical decimal spelling of its validated inputs. That keeps capacity membership, demand caps and expansion comparisons invariant under a supported common rescaling. The full output records the exact fractions and their display approximations. Annual operating cost remains a numerical calculation, and the availability stress remains a declared deterministic change. Exact arithmetic cannot establish achievable capacity, accepted cargo, cost reliability or intervention feasibility.

Decision implications

- Reproduce the complete fixture rather than rounded display values.
- Keep costs and omitted investment obligations explicit.
- Retain demand, tied-limit and zero-service challenges in review.

PRACTICAL APPENDICES

23

Appendix: measurement and intervention records

Field-ready records with explicit missing data

Use records to preserve a decision, not manufacture data

Adapt these completed fictional qualitative records by carrying each named object, actor, decision and missing evidence into a sourced entry for the actual corridor. The examples contain no observed timestamp, incident frequency or empirical operating quantity. Keep unknown values open and preserve the source and status of every replacement. A completed form can still support a conditional decision: it should show which confirmation is missing, who can provide it and what action waits for it. Completion of the form does not mean that the unknown has been measured.

The record should begin with the service edition and the decision it supports. A dispatch record answers whether an eligible object may move into a particular handoff. A measurement record answers how an event will be collected and interpreted. An incident record answers which restrictions and responses apply. A commissioning record answers which requirement has been demonstrated. Keeping these purposes separate prevents a form created for local administration from being reused as evidence of a broader service result.

Each record should distinguish an unresolved property from a failed property and a completed transition. In the fictional dispatch example, missing receipt confirmation leaves destination readiness unknown; it does not establish refusal and does not permit an assumed immediate completion. A confirmed incompatibility is a different disposition and can require a revised formation or destination plan. A confirmed receipt closes that handoff while leaving later vessel acceptance open. Preserving these states prevents an unanswered message from becoming a fabricated operating result. The next action follows the actual state: obtain the confirmation, resolve the incompatibility or advance the cargo trace to its next acceptance event.

The records also retain a challenge. A dispatch approval can be contradicted by missing downstream readiness. A measured interval can be invalidated by incompatible clocks. A claimed buffer can be unavailable because the cargo is unreleased. A commissioning test can fail to represent normal resources. The challenge field is not ceremonial. It directs the review toward the condition most likely to reverse the current interpretation and preserves the reason an earlier decision was qualified.

A dispatch record links eligibility to the next handoff

For the fictional eligible rail formation bound for the port yard and later vessel loading, assign each required confirmation to its originator. The producer supplies cargo identity and release; the rail operator supplies formation readiness and movement conditions; the receiving yard supplies space and acceptance eligibility. The coordinator aligns those confirmations within its agreed authority. Until an actual formation, path and cargo quantity are evidenced, the entry defines a real dispatch's evidence requirement rather than records a movement. Keep each absent confirmation with its owner before deciding whether dispatch can advance.

The record separates planned destination from confirmed receipt. A forecast yard slot can support planning, but the irreversible movement may require an actual confirmation under the relevant operating arrangement. The form should state which confirmation is necessary at the dispatch point and which conditions may remain provisional. If the yard cannot confirm readiness, the response should identify the permitted hold location and affected commitments rather than send the formation into an unmanaged queue.

The cargo record should map the commercial parcel to the physical formation and yard receipt object. A split or combined parcel needs a traceable relationship so that quantity and eligibility can be reconciled. If the mapping is absent, the dispatch may still be physically possible but the end-to-end evidence remains incomplete. The form should not treat a vehicle identifier as sufficient proof that the intended cargo has reached its acceptance boundary.

An exception should record authority, reason and consequence. Dispatch under a temporary recovery priority can be feasible while delaying another user. The record should state who approved the exception and when it expires. A recurring exception should trigger service-basis review. This

prevents the corridor from relying indefinitely on informal interventions that are neither represented in the capacity input nor agreed as part of normal service.

A measurement record protects the observation basis

The measurement example concerns the interval between eligible arrival at the yard and accepted receipt. Its start excludes cargo not yet eligible for the intended service, while a companion record preserves that ineligible waiting. Its end follows acceptance rather than merely the end of physical unloading. This pairing permits both a service-duration measure and a diagnosis of preparation delay. The observation basis should state those distinctions before data are extracted.

The form identifies event time, creation time and transmission time separately. A manually entered event may be accurate in content but late in reporting. A device record may be timely but use an incompatible clock. The collection plan should preserve original values and any correction rule. If a clock discrepancy cannot be resolved, the duration remains unavailable or bounded under a declared method. The example leaves all times unknown because it represents no actual movement.

The observation universe includes eligible arrivals and retains unfinished cases at cutoff. It also records the outside queue and excluded cargo class where relevant to the broader service. This prevents a clean completed-only extract from producing an apparently improved average while backlog remains. The analyst should document the relationship between the local observation and the end-to-end service population, not merely describe the number of available rows.

The cause record distinguishes immediate stop state from attribution. A full receipt area can be observed, while the reason it is full may be uncertain. The collector should retain that difference and seek release, reservation and downstream records before assigning blame. The resulting evidence can support a focused trial or operating change without making an unsupported claim about an actor's motive or contractual performance.

An incident record governs restriction and restoration

The fictional incident concerns loss of inland replenishment while a vessel-loading commitment remains active. The record first checks whether eligible reserve stock can protect that commitment. It identifies

the relevant stop authority, stock eligibility, reclaim readiness and marine permission. If those conditions are unconfirmed, the reserve remains conditional. The form does not assign an interruption duration or claim that the protected service can continue for a quantified period.

The response should specify upstream and downstream actions together. The producer may need to restrain release; the railway may need to protect resources and revise paths; the yard may need to preserve space and eligible stock; the loading actor may need to revise the vessel plan. These actions depend on their respective authority and evidence. The coordinator's role is to align feasible responses and make user impacts visible, not to substitute a system objective for a technical or public stop condition.

The restoration section distinguishes repaired capability, operational readiness, resumed movement and stable service. It also asks whether reserves have been rebuilt. The incident should not close merely because the first train moves or the loader restarts. If temporary priorities and restrictions remain, the service state should retain them. The record should identify the condition and actor that permit return to the normal operating basis.

A corrective action should address the mechanism supported by evidence. If replenishment failed because of a path restriction, a loader-rate improvement is not a relevant closure action. If the yard's reserve was physically present but ineligible, a stock-volume increase may not solve the duty. The incident record should preserve the causal hypothesis, competing explanation and next test. This links operational learning to the intervention sequence rather than treating incidents as isolated stories.

A handover record keeps claims within their evidence

The handover example concerns a proposed intervention after a local component test. The test may establish a defined equipment capability while the recurring corridor service remains unproven. The record should identify the accepted requirement, test conditions, temporary support and residual defects. It should then state the service state that may legitimately begin. This prevents a supplier milestone from being silently promoted into full corridor acceptance.

Update the model basis only as far as the commissioning evidence supports. For an effective annual capacity, retain the observation basis and reproduce the revised fixture through the canonical engine. Where annualisation is unsupported, leave the input provisional and name the

evidence still needed. A required annual cell in a production table supplies no reason to invent capability. The handover should carry that unknown into its claim and next test, keeping the demonstrated local property separate from the annual service result.

An intervention evaluation record should compare the intended mechanism with the observed response under a defined basis. It should record concurrent changes in demand, cargo mix and operating envelope. A shorter local wait may be useful even when attribution is uncertain, but the report should not declare a causal effect without an adequate design. The evaluation should identify the next observation or trial needed to distinguish competing explanations.

Follow the fictional exception from dispatch to handover. Without compatible yard readiness confirmation, the formation stays at the last permissible hold point. The measurement record carries the missing event as unresolved rather than assigns an interval of zero duration. If that hold threatens loading duty, the incident record identifies eligible reserve and its closing condition. A later local equipment test closes its technical defect while destination acceptance remains open at handover. Each form therefore has a distinct job, and the unresolved handoff retains its consequence through the chain. This is the proposed record logic; no movement or test is asserted to have occurred.

Completed proposed dispatch record

Fictional qualitative example. All physical times, quantities and rights remain unobserved or unconfirmed.

Completed proposed dispatch record			
Field	Completed example	Status	Action
Decision	May eligible rail formation advance to port yard?	Proposed service decision	Confirm required readiness at irreversible dispatch point
Cargo mapping	Producer parcel linked to formation and receipt identity	Mapping structure defined; no actual object	Obtain actual identifiers and reconciliation
Receiving readiness	Yard space and cargo eligibility confirmation required	Unknown	Receiving owner supplies authorised confirmation
Movement basis	Compatible path and formation required	Unknown	Rail actor supplies relevant operating evidence
Exception	Recovery priority must name authority and expiry	No exception executed	Record impacts on other users before approval

Completed proposed measurement record

No event has been observed. The record defines collection and interpretation so that missing values cannot be mistaken for instant movement.

Completed proposed measurement record			
Field	Completed basis	Current value	Challenge
Start event	Eligible arrival at declared yard boundary	Unknown	Outside and ineligible waiting needs companion record
End event	Accepted receipt under cargo and quantity basis	Unknown	Unloading completion may precede acceptance
Clock basis	Original event time plus recorded correction rule	Not reconciled	Device, manual and transmission clocks can differ
Universe	Eligible arrivals with unfinished cases retained	No actual cohort	Completed-only extraction can bias interpretation
Cause	Immediate stop condition separate from attribution	Not observed	Do not assign motive from elapsed time

Completed proposed incident and recovery record

The example names a continuity duty without inventing duration or buffer size.

Completed proposed incident and recovery record			
Field	Completed example	Evidence status	Closure condition
Interrupted duty	Inland replenishment unavailable	Fictional scenario only	Relevant technical and operating cause established
Protected duty	Continue permitted loading with eligible reserve	Conditional	Cargo release, reclaim and marine readiness confirmed
Response	Align upstream release, transport and yard acceptance	Proposed	Each actor approves actions within its authority
Restart	Verify capability, access and downstream readiness	Untested	Resumed movement under recorded restrictions
Recovery	Restore stable service and reserve duty	Untested	Repeated accepted handoffs and replenished reserve

Field-record use procedure

Choose the record that matches the decision, fill evidence and unknowns separately, and link it to the current service edition. Preserve the challenged condition and the owner of the next action. Do not convert a proposed record into an observed result or an executed agreement during export or handover.

1. Name the decision object and boundary.
2. Record the source and status of each entry.
3. Retain unknown values and evidence actions.
4. Link subsequent changes to the earlier decision and its conditions.

Completed examples are not completed operations

Use these authored qualitative forms to design collection and governance, then obtain the relevant operating records, authority and independent review before drawing an operational conclusion. No actual movement, test or agreement stands behind the examples. A completed illustrative entry defines a collection or decision duty; its status must be replaced with evidence of the actual event before it is used as a service finding.

Decision implications

- Preserve a decision and its evidence status in every form.
- Keep proposed, observed and accepted states separate.
- Link field records through the service edition and change rationale.

PRACTICAL APPENDICES

24

Appendix: evidence scope and a field mechanism challenge

Specify the observation that would change the intervention before approving its causal justification.

Evidence has an object, edition and scope

The evidence architecture distinguishes institutional guidance, fictional model inputs, reproduced outputs and original analysis. Guidance supports a narrow proposition about a method or process within its original scope. Fictional inputs support an illustration. Reproduced outputs support the mathematical consequence of those inputs. Original

analysis connects mechanisms, challenges and decision frameworks. None should be promoted into another evidence class simply because the paper is attractively rendered or mechanically long.

A source card identifies the public publisher, retained edition or accessed page, locator, supported proposition and limits. Retrieval evidence establishes that readable primary material was obtained and retained. It does not automatically establish that every sentence in the manuscript is supported. The author therefore keeps source-derived notes narrow and separates them from the original system argument. Independent review should inspect the precise locator and the consequence drawn from it, not merely count bibliography entries.

The evidence cutoff defines what this edition had available. A live page can change after retrieval, while an archived PDF preserves a particular document object. A later source revision should trigger a support review if it affects a consequential claim. The paper should retain the earlier receipt and explain the change rather than silently replace the evidence beneath unchanged text. Edition transparency concerns actual source and model changes, not an invented history of reviews that never occurred.

The study defines no empirical operating universe. It therefore asserts no observed bottleneck distribution, cargo failure frequency, regional climate probability or market tariff. Its public institutional sources supply methodological context, not measured parameters for the fictional chain. This exclusion protects the reader from an easy but unsupported transition: using a credible publisher's broad statement to make a precise local capacity claim that the publisher never established.

Historical logistics assessment needs a methodology bridge

The historical World Bank logistics report and its redesigned indicator page have different evidence bases. A perception-based assessment summarises a defined respondent population and methodology. Operational shipment indicators summarise another kind of observation. A change between their displayed results cannot be interpreted as an unchanged time-series movement without a bridge. The paper imports neither score nor ranking; it uses the contrast to explain the need for evidence-object and edition definitions.

A practitioner should ask what the external measure actually observes, which period and population it covers, how missingness is treated and what decision it supports. A national context indicator can be useful for framing diligence, while a corridor movement record can test a local

mechanism. They can inform the same discussion without becoming interchangeable. The bridge should state the relationship and limit of each source rather than combine their authority into an unsupported numerical conclusion.

The same principle applies to maritime outlook evidence. The retained UN Trade and Development source excerpt supports qualitative uncertainty and disruption mechanisms only. Direct full-report retrieval was unavailable in the underlying receipt, and no forecast quantity or market statistic is asserted here. This limitation remains visible. A later successful full-report retrieval would permit additional locator review, but would not retroactively validate a claim that this edition did not make.

Use a broad indicator or outlook to identify the local service question, then collect operating evidence capable of quantifying that service. Declare those inputs before the annual app computes their consequences. An external report cannot validate them through citation, and respected institutional context is not an empirical baseline for an unnamed corridor. Keeping the evidence bridge explicit lets the reader trace which source framed the question, which local record supports the input and which conditional result the calculation supplies.

Evidence and scope

The retained primary excerpt discusses policy uncertainty and disruption in maritime trade. Full-report retrieval was unavailable, so this study asserts no market statistic, forecast quantity or full-report methodology.

Evidence: Review of Maritime Transport

From a capacity claim to a discriminating field packet

An expansion sponsor can point to a congested loader, a larger design rating and the importance of port access without demonstrating that loading governs annual delivered flow. The controlled cases below retain the declared serial service and expose the distinction. They ask which capacity change can alter the governing minimum when every other stage remains on its stated operating basis. That question is narrower than whether the port intervention is operationally valuable, financially attractive or necessary for a particular vessel window. Its answer gives the team a mechanism to investigate before committing to a causal explanation.

Illustrative calculation. Assumptions and sensitivity are published with the model.

Controlled annual chain contrasts

Illustrative canonical outputs under fixed operating-cost assumptions. These annual serial contrasts do not forecast a dated service or appraise the interventions.

Controlled annual chain contrasts			
Frozen case	Delivered Mt/year	Binding stage	Operating USD/t
Original serial chain	11.90	Rail corridor	23.73
Loading intervention only	11.90	Rail corridor	23.73
Rail intervention only	12.00	Vessel loading	23.67
Rail and loading interventions together	14.40	Port stockyard	22.39

The loading-only case retains the original annual flow and rail constraint. The rail-only case moves the binding stage to vessel loading and gains little annual flow because loading then governs the minimum. The coordinated case reaches the stockyard constraint. These contrasts do more than name the initial bottleneck: they expose how an apparently large local expansion can produce a small system change and identify the next constraint that a follow-on proposal must address. Because the operating costs were deliberately held fixed, the lower unit cost in a higher-flow case is an allocation consequence. It is not a verified efficiency saving or return on investment.

A measured shortfall at rail would still require a causal investigation. Intrinsic capacity can be low because of traction, train length, track occupancy or infrastructure condition. Observed flow can also be low because the mine releases irregularly, documents are incomplete, the stockyard cannot receive the train or loading operations make acceptance unavailable. Those are different interventions. A rail utilisation ratio based only on annual effective capacity cannot distinguish them. The field record must classify productive work, intrinsic loss, starvation and blocking with consistent timestamps and reasons, preserving uncertain classifications for review instead of treating all idle time as rail failure.

The observation contract should define the movement object before teams compare records. Identify a consignment or train across mine release, dispatch, receiving acceptance and onward loading. Record the event definition, local time basis, source system and any correction history. Link documents and quality acceptance to the same movement where the operational question requires them. If one team measures gate departure

and another measures accepted receipt, their durations can be combined only after the boundary is reconciled. A missing link should remain a missing link; it should not be replaced with a convenient average that makes the chain appear continuous.

The accepted-service denominator also needs attention. A chain may move physical tonnes that are not yet available for a vessel because quality, documentation or receiving conditions remain unresolved. A port operation may receive cargo earlier and therefore improve service reliability without increasing annual tonnes. The annual model measures its declared physical throughput and costs; it does not silently value documentary readiness or schedule reliability. The committee should state whether it is authorising greater annual flow, more reliable dated service, lower operating cost or a combination, then retain separate evidence for each promised outcome.

A useful controlled operating challenge changes the hypothesised binding mechanism while observing the surrounding interfaces. For example, an agreed dispatch trial could alter the release discipline while keeping the receiving procedure stable. The record would track whether idle rail time shifts from starvation to productive work and whether queues or refusals move downstream. A maintenance intervention would need a different protocol that distinguishes restored intrinsic capability from more favourable arrivals. Neither trial is a causal proof merely because flow increases afterward. The interpretation should account for cargo mix, demand, weather, authority changes and other concurrent influences relevant to the declared service.

The decision response is a staged authority packet. First establish the service and a coherent baseline. Then identify which evidence could distinguish the plausible mechanisms and who can supply it. Use the annual contrasts to bound the structural question and a suitable dated model where windows, variability or inventory govern the intended benefit. A staged trial can precede an irreversible capacity commitment when its result would change the design or scope. Keep the option to stop visible, including the cost and operational burden of further evidence. Repeated analysis without an actionable distinction between outcomes does not justify deferring a decision indefinitely.

A successful intervention should carry an explicit duty to examine the next constraint. Increasing productive rail capability while leaving loading unchanged moves the governing annual stage rather than abolishing the serial minimum. Joint rail and loading capability can then expose stockyard receiving or handling. The operating team should prepare acceptance rules, dispatch interfaces and feasible contingency

arrangements for that transition. A later constraint is not necessarily evidence that the earlier intervention failed; it can be the expected consequence of relieving the earlier limit. The sponsor should distinguish this planned progression from an unexpected interface failure that invalidates the assumed service.

Separate the mechanism from the visible symptom

Original completed diagnostic packet for the fictional annual contrasts. The observation requests are a proposed field design, not a claim that measurements have been collected.

Separate the mechanism from the visible symptom		
Hypothesis	Discriminating operating evidence	Consequence for the decision
Rail has insufficient intrinsic productive capability.	Productive train cycles and cause-coded intrinsic losses under an agreed service boundary.	Test feasible rail capability before valuing a design-rating expansion.
Rail is starved by irregular release.	Linked cargo-ready, document-ready and dispatch timestamps with unmet dispatch opportunities.	Examine release discipline or upstream reliability before adding rail hardware.
Rail is blocked by receiving acceptance.	Receiving refusals, stockyard state and train-ready dwell on the same movement objects.	Repair receiving interfaces and evaluate whether capacity is usable.
Loading affects dated service but not annual flow.	Cargo-ready stock, actual vessel windows and accepted loading events with an explicit calendar.	Use a time-based service model rather than infer the benefit from annual tonnes.
Higher flow appears to lower cost.	Separate fixed allocation, variable consumption and new intervention costs.	Build a full appraisal before calling the denominator effect an economic saving.
Relieving rail moves the constraint.	Downstream capability, acceptance and contingency duties under the changed service.	Commit the next-interface responsibilities with the first intervention.

Unchanged annual flow does not exhaust operational value

A loading intervention can be justified even when the annual serial case shows unchanged throughput. A faster operation could improve a particular vessel window, reduce a contractual exposure or provide redundancy. Those are plausible mechanisms, not results calculated here. The correct response is to formulate the dated benefit, retrieve the relevant evidence and choose a model that can represent it. Rejecting the annual-flow claim does not by itself reject the intervention. It separates the unsupported justification from the justification that still deserves testing.

Carry the mechanism test, rebuild the local evidence

The controlled studies have no empirical operating population, fitted queue distribution or capital-cost estimate. The field packet supplies a way to discriminate mechanisms; it has not been executed. Local applicability depends on cargo, calendar, authority and acceptance boundaries that must be rebuilt for an actual service. The annual app remains a transparent serial-flow illustration, while dated reliability and investment value require additional evidence and suitable methods.

Decision implications

- An annual serial contrast can reject a flow justification without exhausting an intervention's operational value.
- Cause-coded movement evidence distinguishes intrinsic loss from starvation and blocking.
- Relieving a constraint carries a duty to prepare the next receiving and service interface.

PRIMARY EVIDENCE

References and source scope

wb-risk-identification

World Bank PPP Resource Center. Identifying Risks. live guidance, accessed 2026-10-10

Locator: Economic or financial; legal or regulatory

Scope: A list of possible risks supplies no empirical event frequency or project-specific probability.

<https://ppp.worldbank.org/identifying-risks>

wb-corridor-toolkit

World Bank. Trade and Transport Corridor Management Toolkit. 2014

Locator: Corridor performance; corridor components and institutions

Scope: Corridor context and measurement basis must be established locally; toolkit examples are not evidence for a new corridor.

<https://documents1.worldbank.org/curated/en/719971468325781473/pdf/Trade-and-transport-corridor-management-toolkit.pdf>

wb-lpi-2023

World Bank. Connecting to Compete 2023. 2023

Locator: Methodology and logistics dimensions

Scope: Do not equate perception scores with site capacity. This edition predates the redesigned LPI 2.0.

<https://documents1.worldbank.org/curated/en/099042123145531599/pdf/P17146804a6a570ac0a4f80895e320dda1e.pdf>

wb-lpi-current

World Bank. About Logistics Performance Indicators 2.0. 2025 methodology, current page accessed 2026-10-10

Locator: Methodology change

Scope: Edition comparisons require a bridge between methods. No unchanged time-series interpretation is assumed.

<https://lpi.worldbank.org/en/about>

unctad-rmt

UN Trade and Development. Review of Maritime Transport. 2025 edition listed, accessed 2026-10-10

Locator: Chapter on international maritime trade: policy uncertainty and continued disruption

Scope: Direct full-report opening was unavailable. Use is limited to the exact cached primary excerpt; no market statistics, forecast quantity or full-report methodology is asserted.

https://unctad.org/system/files/official-document/rmt2025_en.pdf

green-book

HM Treasury. The Green Book. 5 February 2026

Locator: Options appraisal; risk and uncertainty; monitoring and evaluation

Scope: Its social appraisal perspective and public discount conventions do not automatically apply to private finance.

<https://www.gov.uk/government/publications/the-green-book--2/the-green-book-2026>

magenta-book

HM Treasury. Magenta Book: Central Government guidance on evaluation. 15 May 2026

Locator: Evaluation planning; test and learn; benefits management

Scope: A monitoring improvement is not by itself a causal impact estimate. Identification and context remain necessary.

<https://www.gov.uk/government/publications/the-magenta-book/magenta-book-central-government-guidance-on-evaluation-html>

aqua-book

UK Government Analysis Function. The AQUA Book. current edition accessed 2026-10-10

Locator: Quality assurance of analysis

Scope: A quality framework is not a certification of a particular model. Evidence of executed checks is still required.

<https://www.gov.uk/guidance/the-aqua-book>

orange-book

HM Treasury. The Orange Book: Management of Risk, Principles and Concepts. 29 July 2026

Locator: Part I principles; Part II risk control framework

Scope: UK central government requirements are described in their scope. They are not asserted as duties for every organisation.

<https://www.gov.uk/government/publications/orange-book/the-orange-book-management-of-risk-principles-and-concepts>

nasa-se

NASA. NASA Systems Engineering Handbook. NASA/SP-2016-6105 Rev 2, 2016

Locator: Section 6.8 Decision Analysis; technical management processes

Scope: NASA engineering practice is a reference architecture. Transfer to a commercial organisation requires adaptation.

https://www.nasa.gov/wp-content/uploads/2018/09/nasa_systems_engineering_handbook_0.pdf

gao-schedule

US Government Accountability Office. Schedule Assessment Guide: Best Practices for Project Schedules. GAO-16-89G, 22 December 2015

Locator: Ten schedule best practices; schedule risk

Scope: A deterministic critical path is a model result. It is not a probabilistic completion guarantee.

<https://www.gao.gov/products/gao-16-89g>

gao-cost

US Government Accountability Office. Cost Estimating and Assessment Guide. GAO-20-195G, 12 March 2020

Locator: Technical baseline; assumptions; sensitivity; update with actuals

Scope: An earned-value metric cannot repair an invalid baseline or prove the quality of completed work.

<https://www.gao.gov/products/gao-20-195g>

wb-lifecycle

World Bank. Raising economic efficiency in view of life-cycle costs. 22 April 2024

Locator: Quality Infrastructure Investment principle two

Scope: A lifecycle principle supplies no supplier-specific cost data or performance warranty.

<https://www.worldbank.org/en/brief/2024/04/22/qii-2-raising-economic-efficiency-in-view-of-life-cycle-costs1>

wco-trs

World Customs Organization. Time Release Study Guide. Version four, current page accessed October 2026

Locator: Introduction and methodology enhancements

Scope: Institutional methodology supplies no release-time observation for the illustrative corridor.

<https://www.wcoomd.org/en/topics/facilitation/instrument-and-tools/tools/time-release-study.aspx?p=1>

imo-jit

International Maritime Organization. Just in Time Portal. Current portal accessed October 2026

Locator: Implementation; data categories; key timestamps

Scope: No emissions percentage or port-specific effect is imported; actual rights and service feasibility require local review.

<https://greenvoyage2050.imo.org/just-in-time-arrivals/>

imo-facilitation

International Maritime Organization. Facilitation. Current page accessed October 2026

Locator: Government controls and facilitation overview

Scope: The overview is not a legal opinion on a particular port or proof of local implementation.

<https://www.imo.org/en/ourwork/facilitation/pages/default.aspx>

wto-tfa

World Trade Organization. Agreement on Trade Facilitation. Legal text accessed October 2026

Locator: Article seven, pre-arrival processing; Article eight, border agency cooperation

Scope: Applicability, member commitments and local implementation are jurisdiction-specific; no entitlement for an unnamed shipment is asserted.

https://www.wto.org/english/docs_e/legal_e/tfa_e.htm

fhwa-reliability

US Federal Highway Administration. Application of Travel Time Data and Statistics to Travel Time Reliability Analyses. Handbook chapter one, accessed October 2026

Locator: Trip-Based Travel Time Reliability; Standard Processing Procedures

Scope: US road-method guidance transfers only as a measurement concept; no local freight variability estimate is assumed.

<https://ops.fhwa.dot.gov/publications/fhwahop21058/chap1.htm>

wb-port-governance

World Bank. Port Reform Toolkit: Port Governance. Third edition, Module 3; retained 2025 publisher PDF, retrieved 2026-10-10

Locator: Sections on port authorities as community builders and cooperation beyond local jurisdiction

Scope: The public module supports a governance architecture, not an executed arrangement or a local assignment of legal powers for this study.

<https://documents1.worldbank.org/curated/en/099073025114030998/pdf/P176624-d39fa407-886c-4542-96ec-9af4971eb97a.pdf>

networkrail-capacity

Network Rail Infrastructure Limited. Network Statement for the timetable year. 2026 Network Statement, retrieved October 2026

Locator: General information and structure; capacity allocation; referenced engineering access and timetable planning rules

Scope: The Great Britain timetable and access framework is a scoped example, not a universal legal or technical requirement. No local train-path capacity is inferred.

<https://www.networkrail.co.uk/wp-content/uploads/2024/11/2026-Network-Statement.pdf>

PUBLICATION RECORD

Methods, limitations and review

This edition uses publicly accessible primary sources. All worked numerical examples are illustrative, with frozen inputs and reproducible calculations. They represent no actual client, mandate or operating project.

Publication fields identify the fixed edition dated 2026-10-10. Independent assessment and release verification are recorded against that fixed edition separately; these fields do not assign a quality score.

2026-10-10: Initial substantive authored draft replaces the production outline. The study contains explicit corridor mechanisms, fictional canonical input/output studies, proposed practitioner records and original conceptual diagrams. Mechanical and independent acceptance remain separate recorded states.

2026-10-10: Added explicitly fictional computational boundary studies and revised canonical model. I.O precision policy. Independent assessment remains pending.