How Southeast Asia Logistics Gain From Aggregating AI Opportunities

July 2026

By Vincent Chin, Johannes Distler, Lennon Loy, Faye Sit

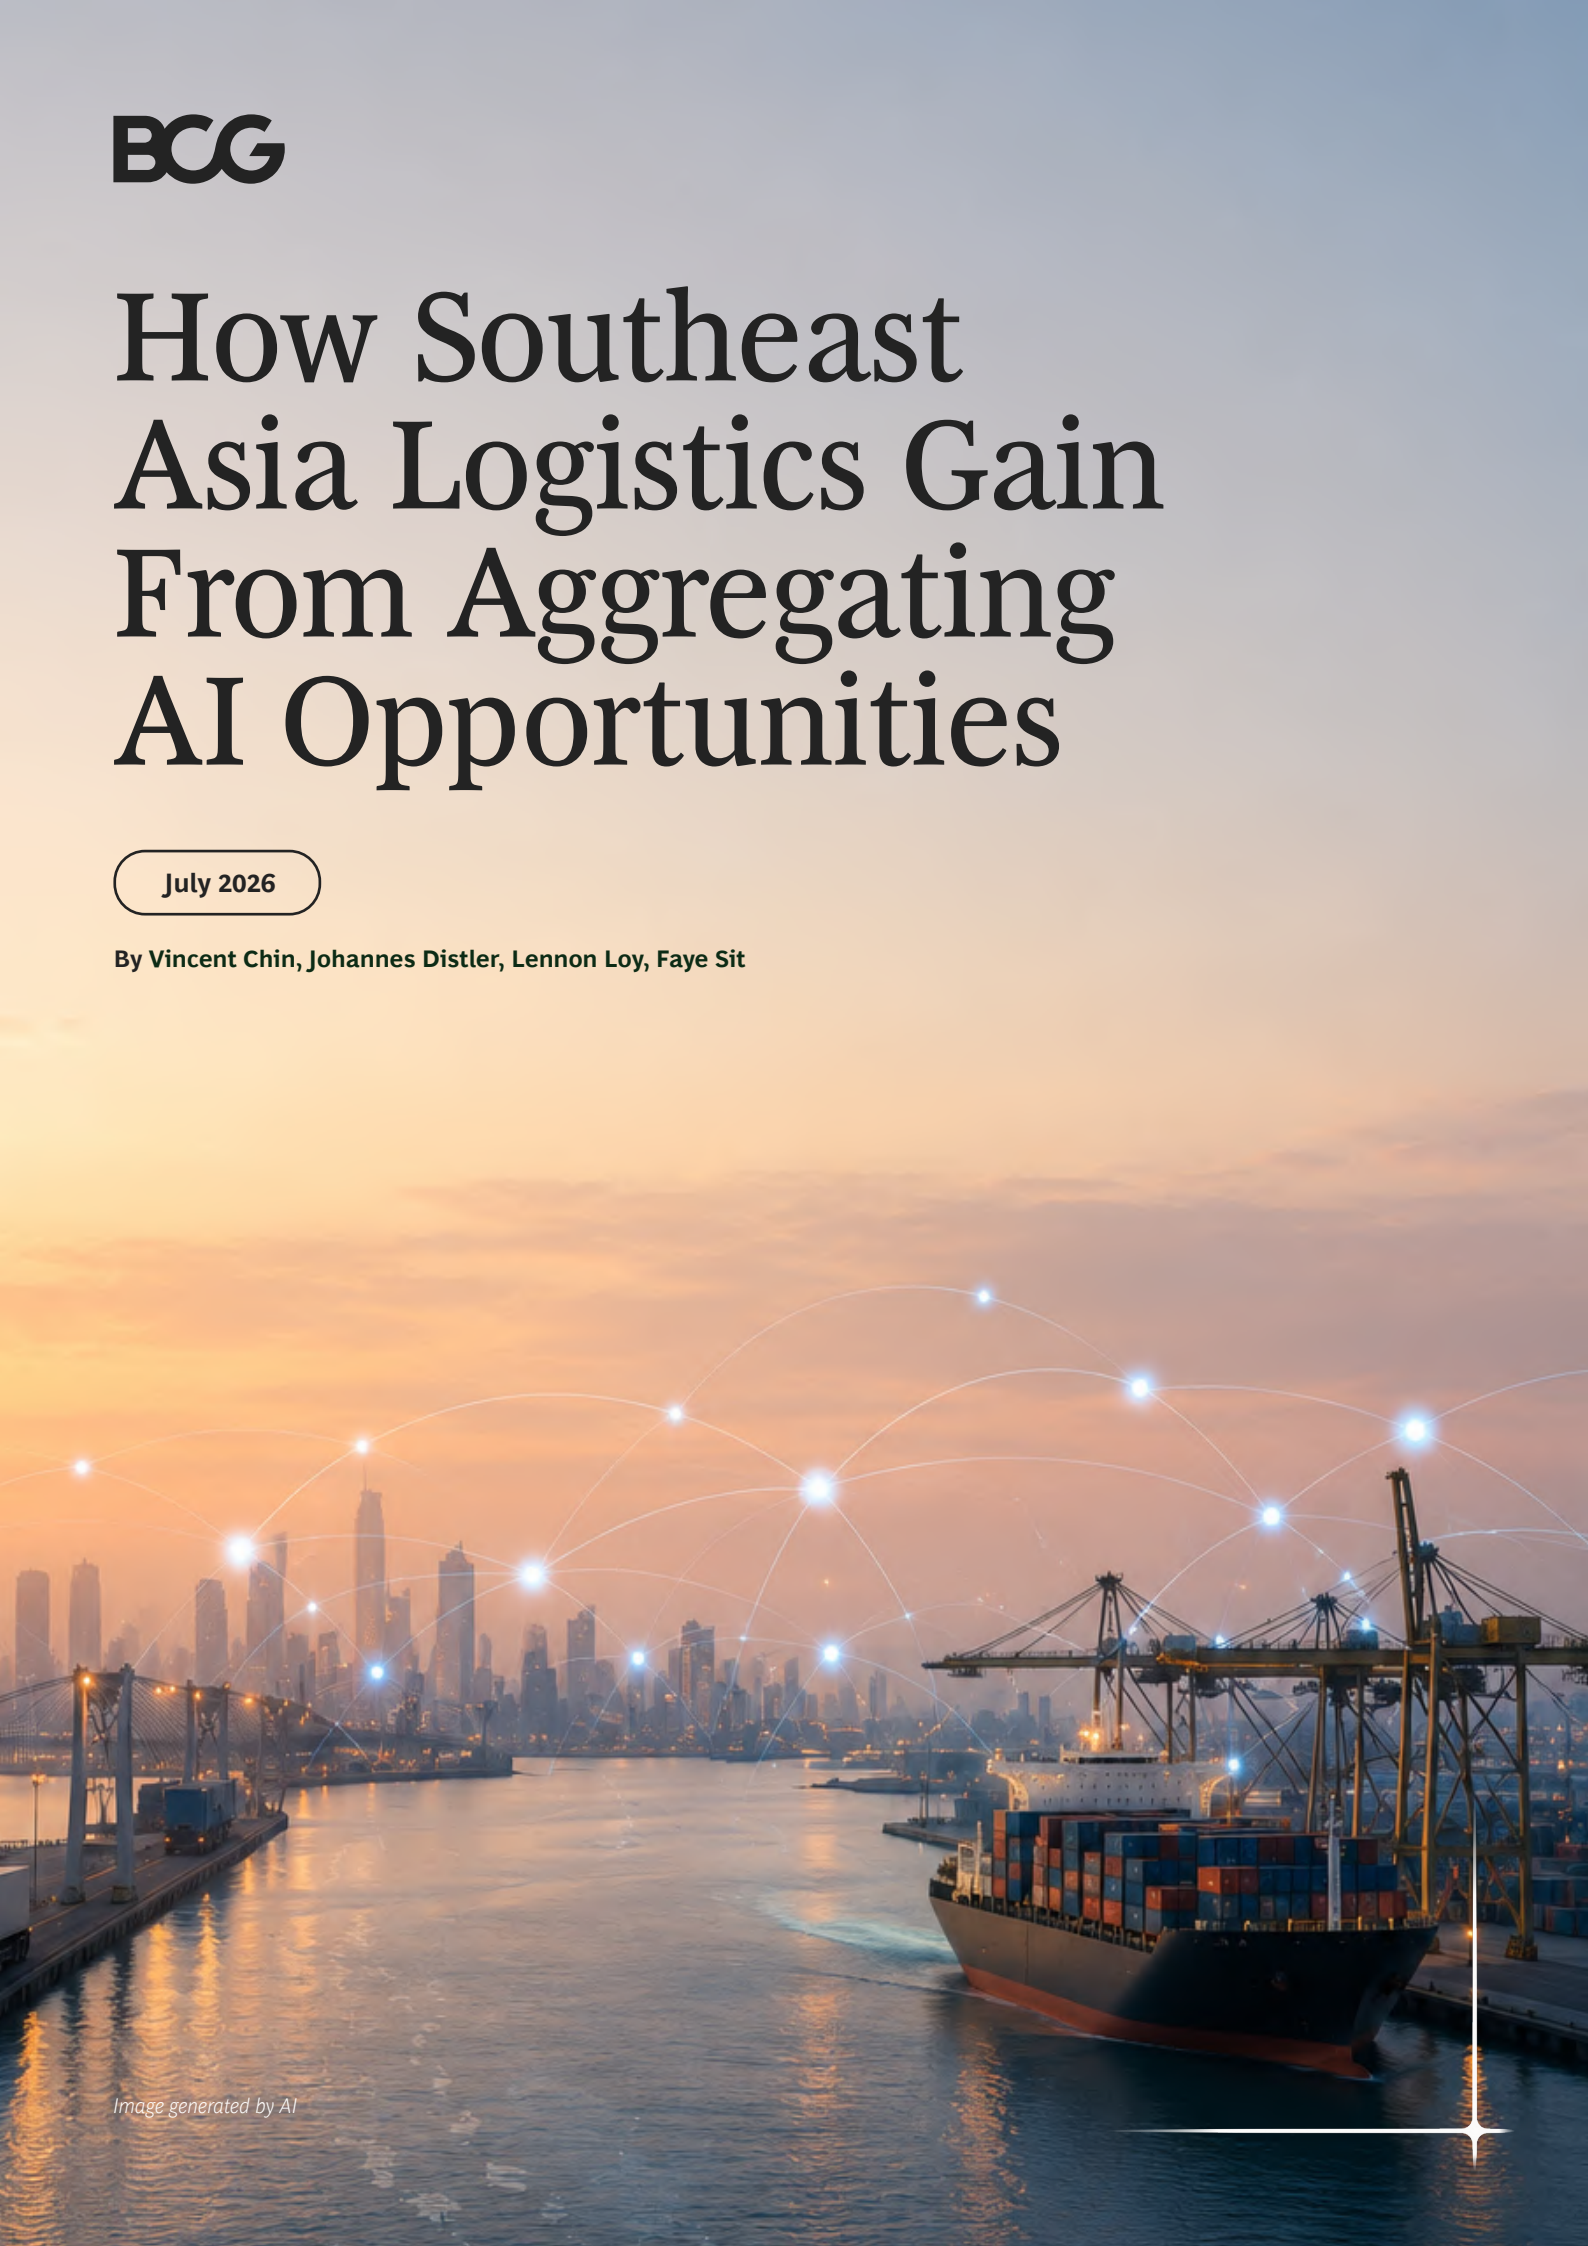

Image generated by AI

Executive Summary

Capture value created by AI initiatives from internal operations and other links in the logistics chain

Logistics AI spending is rising, but returns often fall short. The cause of this problem is structural. Companies deploy AI for isolated tasks. Improvements are often limited to individual points rather than linked up across partners in the wider chain.

1. Adoption has increased faster than measured returns

Only 13% of logistics service providers (LSPs) report measurable financial impact from AI and 7% of shippers report attributable supply chain improvement. This is despite cost reduction motivating close to 80% of investment.

2. Point-solutions constrain deployment, not capability

Most logistics AI is deployed within a single company's boundaries. Gains in precision generated upstream are absorbed into the next player's manual process so the impact remains limited.

3. Chained improvements can compound

When partners structure data on their activities to be easily used and built on by the next player, variance reduced at one point in the chain lowers the uncertainty the next must plan around. Ultimately, these connected upgrades create a total impact far greater than isolated fixes.

4. Internal capability and external linkage are strongest together

Acting alone, an operator can already gain far more by fusing its own data with open-source information about the rest of the ecosystem than by optimizing its operations in isolation, with no partner required. External linkage then compounds that into a position that single-layer competitors find much harder to replicate.

5. Conditions in Southeast Asia are favorable

Asia-Pacific (APAC) leads global logistics AI adoption at 31% of LSPs, against 14% in North America and 6% in Europe. Combined with hub-and-spoke trade and an SME-dominated landscape in Southeast Asia, the region presents a combination of conditions less common elsewhere. Southeast Asia reflects a focused part of this wider APAC opportunity with significant potential for regional players.

6. A new front where mid-sized players can outdo larger rivals

Chain-linking has opened a line of differentiation that does not depend on scale. Because the advantage comes from being first to publish data the rest of the chain can build on, a focused, mid-sized operator can outdo larger, more vertically integrated or more established players that move later, setting the standard and format others adapt to.

The Adoption Gap

Returns from AI investments have not always paid off for LSPs and shippers

AI is firmly on the logistics agenda. Yet by the industry's own account, the financial returns measured so far have been modest, and closing the gap between ambition and outcome is now the central challenge of adoption.

Boston Consulting Group's (BCG) January 2026 survey of over 180 LSPs and shippers across four regions finds that 13% of LSPs report measurable financial impact from AI. [Exhibit 1] Among shippers, 7% report supply chain improvements they attribute to AI investment. Close to 80% of respondents cite cost reduction as the primary motivation for that investment, which makes the shortfall in measured return notable.

EXHIBIT 1

13%

LSPs reporting measurable AI impact[1]

7%

Shippers reporting supply chain improvement attributed directly to AI

~80%

Citing cost reduction as primary goal the dominant investment rationale

The gap is not closing

It is tempting to assume that AI returns are simply delayed, and that profits will follow as teams gain experience. However, our survey data complicates that reading. Adoption has accelerated faster than measured impact, so the distance between the two has widened rather than narrowed. If the relationship were a simple time lag, the leading adopters would be converting investment into return at materially higher rates. However, the data does not show this.

This suggests a constraint that further spending and better models are unlikely to resolve on their own. The relevant question is not how much AI is deployed, but how the value it creates moves, or fails to move, through the sequence of players that handle a shipment from origin to destination.

THE SO WHAT

If the gap between AI investment and AI impact is structural and not closing, doing more of the same is unlikely to help. The question worth asking is how to deploy AI investment effectively, not simply how much more to spend.

[1]BCG / Alpega survey, Jan 2026, n=180+

The Point-Solution Trap

Why the gains stay small

Most AI in logistics is deployed as a point solution. It is built on a single company's own data, optimized for its own operations in isolation, then handed to a counterparty that absorbs the precision into a manual process.

Let us use the illustrative example of how a single container moves. The carrier optimizes its slot allocation. The port optimizes its berth schedule. The forwarder optimizes its routing and documentation. The trucker optimizes the final leg. Each has a legitimate use of AI, and each delivers a real gain within the boundary of the company that built it.

The constraint starts to appear at the handoff. The carrier's improved arrival estimate reaches the port as an email or electronic data interchange (EDI) message in a planner's queue. The port's refined berth window reaches the forwarder as a portal update checked manually. At each boundary, the precision generated upstream meets a manual process downstream, and much of it is lost. The gain is real but contained: it improves one player's internal metrics without improving the journey end to end.

POINT SOLUTION TODAY

AI output is built for internal use and improves the builder's own KPI. At the handoff it is reduced to a human-readable message that the next player reinterprets manually. Precision decays at each boundary.

CHAIN-LINKED ALTERNATIVE

AI output is built to be read by the next player's system, so precision survives the handoff. Each improvement narrows the uncertainty the next step faces, and gains accumulate along the journey rather than resetting at each boundary.

This is not a shortcoming of the models. The constraint is architectural: AI built for internal consumption cannot compound, because compounding requires the output of one step to become the input of the next. Each has a legitimate use of AI, and each delivers a real gain within the boundary of the company that built it.

THE SO WHAT

AI implementation needs to be set up to support and to maximize the benefits from other links in the logistics chain. For example, the berth-optimization model can perform well internally and still add little end-to-end, if its output is reduced to an email the moment it crosses the fence line.

The Chain-Linking Principle

How small steps combine

A useful reference point comes from outside logistics. It illustrates how a sequence of individually unremarkable steps can produce an outcome that none of them achieves alone.

In a major test of AI capabilities, a large language model (LLM) platform undertook a detailed cybersecurity experiment in 2026. The LLM used its capabilities to gain control of a computer system—not by exploiting a single critical vulnerability, but by identifying four individually unremarkable pieces of information, of the kind a system routinely makes available, and used each to enable the next step.

The first piece revealed where the system's controls were located. The second used that location to find a specific component. The third placed a counterfeit credential at that component. The fourth prompted the system to recognize the credential as an administrator. No single step was an intrusion in isolation, and a security team reviewing any one of them would have found little of concern. Each step supplied what the next required, and in sequence they produced full administrative control.

How precision multiplies

In a connected supply chain, delays cascade to every partner. Reducing disruption early in the journey does more than improve that single stage. It lowers the uncertainty every downstream partner must manage.

Improving the predictability of port timing improves berth planning, which tightens the forwarder's transit commitment, which tightens the shipper's production schedule, which smooths the trucker's gate arrival. Isolated upgrades allow that company to capture only the direct value generated. External partners miss out because they lack visibility and the value of the upgrades are not fully realized. However, connecting these improvements across the entire network multiplies the value for everyone involved.

This explains why total gains easily exceed isolated wins. A standalone tool captures value at one stage. A connected chain captures that initial win and multiplies it across every subsequent handoff. That compounding effect is the true source of value.

THE SO WHAT

The parallel for logistics is that the value may not sit within any single player's operations. It may sit in the sequence, with each improvement supplying the input the next improvement needs.

Why Southeast Asia

Conditions suited to compounding

Several structural features make the chain-linking argument more applicable in Southeast Asia than in many comparable regions. Where the region leads is also where it has the most to gain.

APAC leads the world in logistics AI adoption, with 31% of LSPs reporting AI embedded in core operations: more than double North America (14%) and around five times Europe (6%). That position is a starting point, not a conclusion. [Exhibit 2] Three structural reasons explain why, alongside APAC's high AI adoption, Southeast Asia is poised to benefit.

EXHIBIT 2

31%

APAC LSPs with AI in core ops highest of any region surveyed

14%

North America less than half APAC

6%

Europe around one-fifth of APAC

Hub-and-spoke concentration

The region's trade architecture funnels a large share of traffic through a small number of transshipment hubs. The same link often recurs across many journeys, so an improvement made once at a hub applies to a large volume of cargo. A single well-placed improvement carries more leverage than in a dispersed network.

Archipelagic complexity

Indonesia spans roughly 17,500 islands and the Philippines more than 7,600. Multi-leg journeys with sea, port, and road handoffs are the norm. Each additional handoff is both a point where precision is currently lost and a point where chain-linking can recover it.

An SME-dominated landscape

The region's logistics sector is dominated by small and mid-sized operators that lack the scale to build proprietary AI independently. This is usually framed as a barrier. It is also a precondition for the model described here: where counterparties cannot build their own capability, an operator that offers a useful free tool gains both adoption and the data it generates. Crucially, the operator that does so need not be the largest. The advantage goes to whoever moves first in a corridor, which puts a focused mid-sized player on equal footing with larger and more established rivals.

THE SO WHAT

Fragmentation in Southeast Asia is not only a problem to be managed. For an operator that designs for interoperability and moves early, it is also where the opportunity lies, and that opening is open to mid-sized players as much as to the incumbents.

Port Operators

The anchor of the chain

The port sits at the intersection of vessel, cargo, and truck flows. That position makes it the logical anchor of a connected chain, and the operator best placed to benefit from publishing rather than retaining its data.

The analysis that follows takes a deliberately conservative starting point. In this scenario, no counterparty has AI, and many are not fully digital. The only inputs available are open-source data, namely automatic identification system (AIS) vessel positions, public container tracking portals, weather feeds, and published gate schedules. The example illustrates what becomes possible from that baseline, and where modest investment changes the picture.

With open data only

AIS vessel position data is publicly available via MarineTraffic and VesselFinder. Free tiers carry a delay and rate limit. Real-time commercial feeds run from a few hundred to a few thousand US dollars a month depending on coverage and refresh rate, a meaningful sum for a small feeder operator or Tier 3 port but modest relative to the planning value at scale. Ports have historically relied on carrier-reported ETAs, which in practice carry significant buffer, though shipper and forwarder demand for visibility is now pushing some carriers toward greater precision.

A port that builds its own AIS-based arrival model can achieve materially better berth arrival accuracy than one working from carrier figures. The leading hub operators are already well advanced here: PSA Singapore is a recognized benchmark in port digitalization through platforms such as CALISTA and OPTIMAS, and Pelindo has consolidated its Indonesian terminals onto common digital systems. The opportunity described below is therefore most pronounced at Tier 2 and Tier 3 ports, where carrier-reported ETAs remain the primary planning input and AIS-based modeling represents a step change rather than an incremental improvement.

How might an investment in the chain pay back

A port operator that publishes a free berth visibility dashboard for feeder operators creates an asymmetric exchange. The tool costs feeder operators nothing, and in return they confirm arrival via the platform, generating a machine-readable signal more accurate than current inputs. Data quality improves as more operators use it. This pattern is seen in other industries: a free tool for counterparties generates proprietary data for the platform owner. The investment case does not depend on counterparty AI adoption—it depends on counterparties finding the tool useful enough to use.

> Publishing an arrival model as a real-time feed need not mean giving away advantage. The port can become a central link in the network, and the gate data returning through the chain improves its own planning. This is most striking at Tier 2 and Tier 3 ports: an operator that moves first can establish itself as the corridor's reference feed and build a data position that larger, more established hubs would then have to displace.

Freight Forwarders

The intelligence layer

Forwarders touch more counterparties than any other player in the chain. That breadth allows them to aggregate fragmented public data into intelligence no single shipper or carrier can replicate.

With open data only

Every major carrier publishes container tracking on its website. A forwarder that aggregates tracking across all client shipments and all carriers into a single view has something shippers cannot easily build for themselves—a consolidated picture of everything in transit. The secondary value may be larger. Tracking data accumulated across hundreds of shipments and dozens of carriers over time produces a carrier performance database. Consistent late performance on specific corridors, identified through pattern analysis of public tracking data, informs carrier selection and rate negotiation. The advantage lies not in access to the data, which is public, but in aggregating and analyzing it systematically.

How might an investment in the chain pay back

A modest investment in a free tracking application yields outsized returns for forwarders. Shippers adopt it because it consolidates multiple carrier portals into one view. The forwarder receives, in return, consolidated cargo data across all clients: volumes, trade lanes, carrier preferences, and growth patterns. That data describes the client base more accurately than a sales process could. Shippers see an improvement in service—the forwarder builds a commercial intelligence capability from operational data.

Why the forwarder's position is structurally advantaged

A carrier sees its own vessels. A port sees its own calls. A shipper sees its own cargo. The forwarder sees across all of them, because its role is to broker between them. That cross-counterparty visibility is the scarce asset. When a forwarder turns fragmented public tracking into a performance benchmark, it competes not on access but on synthesis, which is harder to replicate and more defensible. That is also why a mid-sized forwarder can outdo a larger one here—the benchmark is built from data already public, so the edge comes from being early and systematic rather than from balance-sheet size.

Carriers and Last-mile

Precision at the edges

The ends of the chain generate the most granular real-time data. Publishing it, rather than consuming it internally, is the step that closes the loop back to the port.

FOR FEEDER CARRIERS

With open data only

Every major carrier publishes container tracking on its website. A forwarder that aggregates tracking across all client shipments and all carriers into a single view has something shippers cannot easily build for themselves—a consolidated picture of everything in transit. The secondary value may be larger. Tracking data accumulated across hundreds of shipments and dozens of carriers over time produces a carrier performance database. Consistent late performance on specific corridors, identified through pattern analysis of public tracking data, informs carrier selection and rate negotiation. The advantage lies not in access to the data, which is public, but in aggregating and analyzing it systematically.

How might an investment in the chain pay back

Among mid-tier and feeder carriers operating in Southeast Asia, it is still common to see ETAs set weeks in advance and revised infrequently. A carrier that publishes rolling 72-hour ETAs driven by live AIS position has a customer service differentiator that is straightforward to build and requires no counterparty investment. Shippers and forwarders making connection and storage decisions on more accurate ETAs are more likely to route cargo through a carrier whose reliability they can trust.

FOR TRUCKERS AND LAST-MILE OPERATORS

With open data only

Port gate congestion at most Southeast Asia ports follows consistent daily patterns. A trucker who logs these patterns and schedules gate arrivals in off-peak windows reduces unproductive waiting without any technology investment. This simple habit demonstrates how AI works: using available data to remove guesswork creates immediate benefits down the line. At this level the barrier is not capability but digitalizing the habit and information-sharing within driver communities.

How might an investment in the chain pay back

Google Maps was not designed for freight. It lacks the freight-specific data that matters on the ground in Southeast Asia: gate hours and congestion patterns, which industrial estates have one-way heavy-vehicle entries, which roads are impassable in monsoon season, and where fuel stops can take large vehicles. Offering drivers a free, truck-specific navigation app creates a powerful exchange of value. Drivers adopt the tool because it immediately improves their daily routes. In return, the logistics company captures real-time GPS, actual route data, and delivery confirmations at no extra cost. As the user base grows, this creates a proprietary traffic network no single operator could build alone.

CASE STUDY: THE TRUCKER NAVIGATION PLATFORM

The Waze pattern, applied to freight

- **The precedent:** Waze gave away consumer navigation. Google acquired it for US$1.3 billion in June 2013, in large part for the traffic data its user base had generated.[2]
- **The compounding effect:** More drivers generate better traffic data, which improves routing, which attracts more drivers, which widens the data advantage. At scale the platform becomes an accurate freight intelligence layer on its corridor. No counterparty needs to invest for this to work. A logistics equivalent has not yet been built at scale in Southeast Asia.

[2]Google acquisition of Waze, US$1.3bn, June 2013. Reported by WSJ, TechCrunch, and the Financial Times at announcement.

The Connected Chain

How the steps connect

The wins described above are powerful on their own. However, the real breakthrough happens when partners format their data to be shared. The following sequence isn't a mandated coordination program. It is a description of the natural advantage that emerges when companies design their systems to work together.

The five-step loop: each output is the next input

From isolated improvements to a connected efficiency layer

PORT OPERATOR

Builds an AIS-based arrival model. Publishes dynamic berth availability through a free dashboard that feeder operators use to confirm arrival. Each confirmation is more accurate than a carrier-reported ETA, and returns to the port as a machine-readable planning signal.

FEEDS STEP 2: Feeder carriers receive accurate berth slot data to schedule against

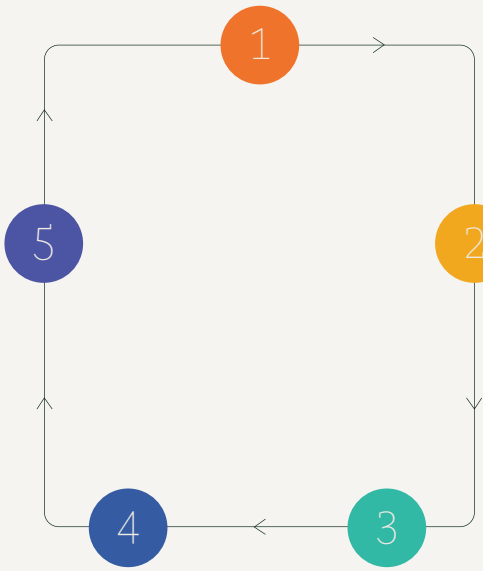

TRUCKER / LAST-MILE

A GPS fleet on the freight navigation platform generates real-time gate congestion data. The port sees, in near-real time, how many trucks are inbound and from where, without requesting reports. As the network grows, intelligence improves, attracting more drivers and widening the advantage.

LOOPS TO STEP 1: Port refines its arrival and gate model with actual vehicle data

FEEDER CARRIER

Reads port berth availability and supplements it with live AIS positions of its own vessels. Publishes rolling 72-hour ETAs driven by data rather than schedule buffers. Precision is now possible because the port's estimate has replaced the carrier's conservative guess.

FEEDS STEP 3: Freight forwarders receive a reliable eta they can commit to shippers

SHIPPER

Aligns factory scheduling and inventory planning against reliable delivery windows. Cargo is ready earlier and gate arrivals are more predictable, reducing the late-gate variability that disrupts port yard sequencing. No technology investment is required. The benefit flows from the reliability upstream.

FEEDS STEP 5: More predictable cargo readiness smooths truck inbound flow to the gate

FREIGHT FORWARDER

Ingests carrier ETAs alongside its own cross-carrier tracking history. Issues a transit commitment to the shipper based on verified carrier performance rather than declared schedules. Its tracking database, built from aggregating public portals, provides a benchmark no shipper could build alone.

FEEDS STEP 4: Shippers receive delivery windows accurate enough to plan production against

THE SO WHAT

Each step's output narrows the uncertainty the next step faces. When you fix errors early, you protect everyone downstream from delays. The total value of these connected improvements far exceeds what any single company could achieve alone.

What the connections require

None of these connections require formal bilateral agreements or a new industry standards body. What they require is a lightweight common format and shared semantics, so that one player's output is not only machine-readable but means the same thing to the player consuming it. An ETA is useful downstream only if both sides agree on what event it marks and in which time zone. Workable de facto formats already exist for most of these signals, so the gap is one of adoption rather than invention.

In practice this means formatting each output for external consumption: machine-readable ETAs rather than schedule PDFs, API-accessible berth status rather than notice boards, standardized tracking events rather than carrier-specific portals. The choice is made at the point of building, costs little extra, and determines whether the gain stays contained or enters the chain.

The Internal Foundation

Internal strength, external reach

The most defensible position pairs strong internal AI capability with external linkage. The first step is fully within an operator's control. By combining its own data with open-source information about the rest of the ecosystem, a company can gain far more than by optimizing its own operations alone, well before any formal partnership is in place.

Chain-linking depends on a counterparty eventually cooperating, which sits partly outside the operator's control. A company can act on its own initiative, enriching its proprietary data with open feeds described earlier: AIS vessel positions, public container tracking, weather and gate schedules. With these, the returns start to accrue whether or not the chain follows.

This is the second of three steps, not the trap described earlier. That trap occurs when a company builds an isolated model using solely internal data for internal tasks. Fusing that data with the open ecosystem is different in kind—the company captures the value by optimizing the entire journey instead of just its internal operations. Publishing that output to the chain is the third step, where the gain compounds as counterparties take up a feed that serves them.

OWN DATA ONLY

A company that optimizes only its own data and processes captures a real but contained gain, of the kind rivals can match by buying the same tools.

OWN DATA PLUS THE OPEN ECOSYSTEM

Fusing proprietary data with open-source information about the wider ecosystem yields far more, and needs no partner to cooperate. Layering external linkage on top compounds it into a position single-layer competitors struggle to dislodge.

THE SO WHAT

The first moves do not require a partner. Combining an operator's own data with open information instantly unlocks more value than optimizing in isolation. For a mid-sized operator that is the accessible starting point, and the linkage that follows turns a solid internal capability into a position larger rivals struggle to match.

The Design Decision

Risks and considerations

The approach set out here carries trade-offs that warrant attention before any operator commits to publishing data externally.

1. Cybersecurity and data ownership

Opening an output as a live feed widens the attack surface and raises the question of who owns the data a counterparty generates by using a free tool. Access controls, clear terms of use, and a defensible ownership position should be settled before launch, not after.

2. Interoperability has a real cost

Formatting output for external consumption is inexpensive at the point of build, but not free. Maintaining an API, versioning a schema, and supporting external consumers is an ongoing commitment that needs a named owner and a budget line.

3. Organizational change

Publishing accurate performance data exposes it internally as well as externally. The harder shift is cultural—moving from buffers and informal commitments to measured, accountable performance. This requires leadership support.

4. Concentration and competition risk

If one player becomes the dominant link in the chain, it gains pricing power, and counterparties feeding it data may find the relationship harder to exit. That position is the strategic prize for a first mover, but it may also attract regulatory and competition scrutiny as it grows.

THE SO WHAT

The case for chain-linking is conditional, not automatic. It holds where an operator can manage the security, cost, and governance implications of publishing data, and a deep understanding of where the inbound data gained justifies the exposure given away.

What to do with this

The practical implication is narrow and concrete. When a logistics player in Southeast Asia scopes its next AI investment, the question is not only what to optimize, but whether the output is built to be read by the next player in the chain. A model built for internal consumption captures a single, isolated gain. The same model, published as a real-time feed, enters a chain that returns better data to its own planning system. The architecture of the output, open versus closed, and machine-readable versus human-readable, is the decision that determines whether individual investments remain point solutions or aggregate into something structurally different.

Where the aggregate return arises

Every player described in this report can start with open data and without counterparty cooperation. The decision that separates point-solution returns from aggregate returns is made at build time, not at scale.

The individual gains are real: better berth prediction, more accurate carrier benchmarking, more efficient truck routing. On their own they are modest. Their wider significance follows the principle illustrated earlier—each improvement generates an output that, if readable by the next player, feeds that player's decision. Connected chains multiply value in ways that isolated upgrades simply cannot.

The barrier is commercial, not technical

Sharing data that reveals operational performance creates competitive exposure. A port that publishes precise berth availability also reveals its own utilization. A carrier that publishes accurate ETAs becomes accountable when it misses them. Players will weigh this carefully, and early movers are likely to be those that judge the inbound data quality they gain to outweigh the transparency they give away. That calculation shifts further toward participation as adoption spreads, which is why the design decision matters most at the point of initial build.

An opening for mid-sized operators

This is also where a mid-sized operator can gain ground on larger rivals, though not automatically. The advantage no longer rests on the deepest pockets or the most vertically integrated network. It rests on being the first in a corridor to publish data the rest of the chain can build on.

A focused, mid-sized port, carrier, or forwarder that moves early sets the de facto format others adapt to, and accumulates the inbound data advantage before larger and more established players commit. Large competitors can quickly copy this technology using deeper pockets. The true protection is the data and user base you built first. This early head start creates a structural advantage that well-funded rivals will struggle to break.

> **For logistics players in Southeast Asia, the question worth asking is not whether to adopt AI, but whether the internal capabilities being built today are also designed for interoperability with the rest of the chain. The operators that pair internal strength with external linkage will be the hardest to displace.**

About the Authors

Vincent Chin
Managing Director and Senior Partner,
Global Vice Chair of Public Sector
Singapore
Chin.Vincent@bcg.com

Johannes Distler
Managing Director and Partner
Dubai
Distler.Johannes@bcg.com

Lennon Loy
Principal
Singapore
Loy.Lennon@bcg.com

Faye Sit
Lead AI Product Manager
Singapore
Sit.Faye@bcg.com

About Boston Consulting Group

Boston Consulting Group bridges the gap between ambition and outcomes for the world's leading companies and organizations. We are built for this era of unprecedented change — bringing strategic clarity rooted in over 60 years of deep domain knowledge, combined with applied AI shaped by our practitioners. BCG works shoulder-to-shoulder with CEOs across industries and geographies to deliver transformative impact at scale: stronger returns, transferred capabilities, and change that sticks. For more information, visit bcg.com.