Last Mile LEADERS
Hosted By FarEye

2026

EYE ON THE LAST MILE 7.0

AMER REPORT

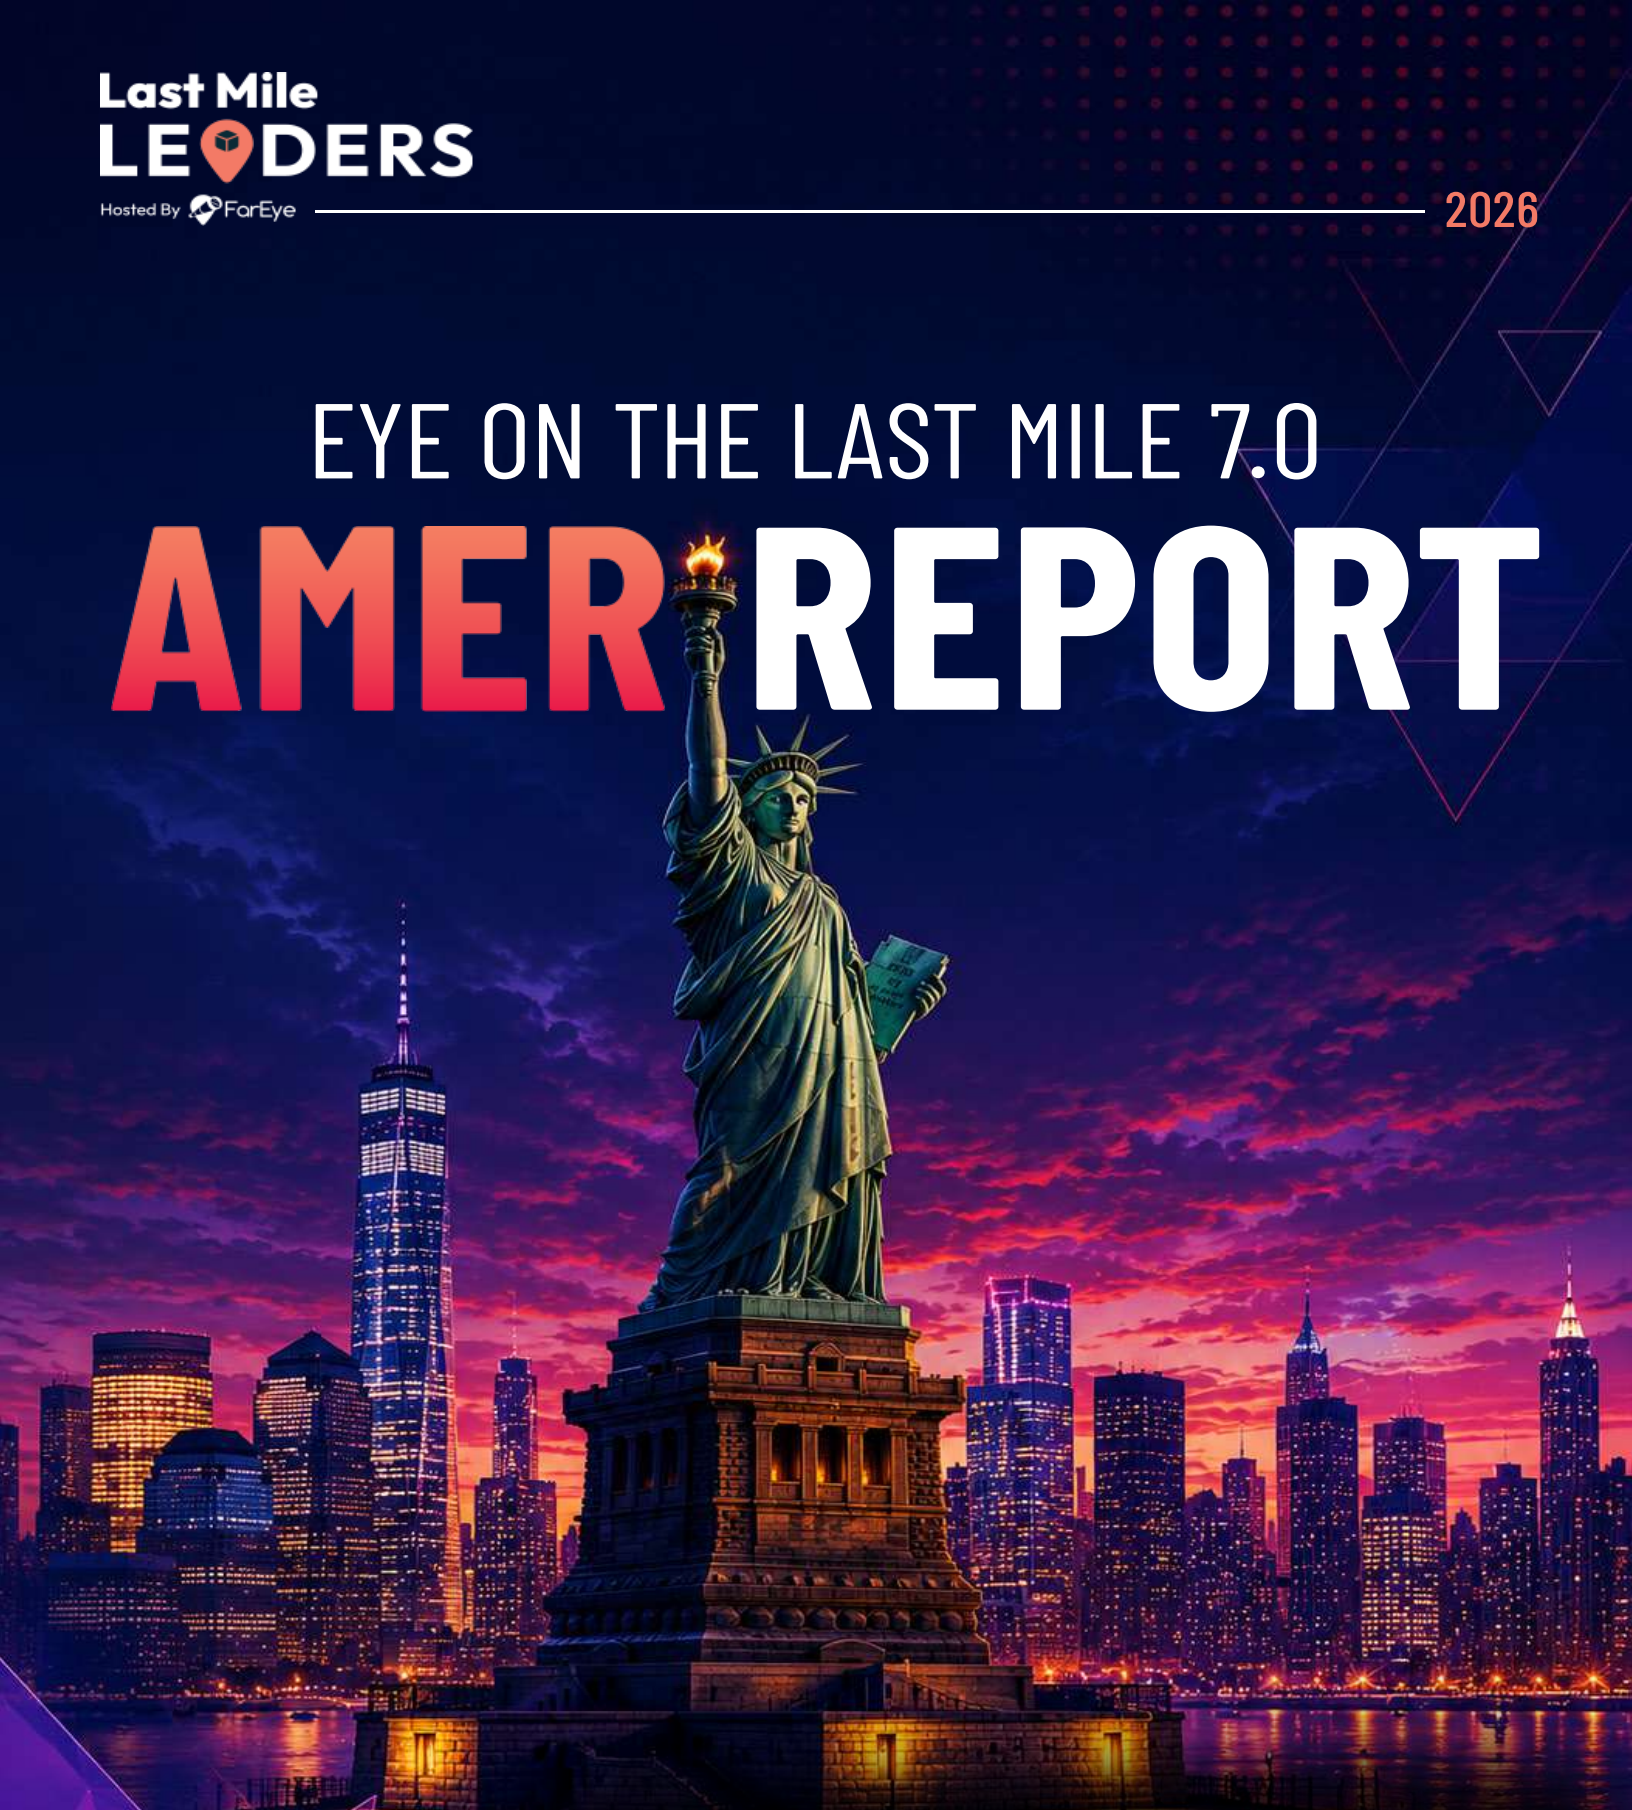

THE LAST-MILE CONVERGENCE: WHAT AMER LAST MILE LEADERS MUST DO NEXT

THE THESIS

Winning the last mile is no longer about maximizing service or minimizing cost. It is about designing a delivery promise customers value, keeping it reliably, and stopping the cost of that promise from scaling as fast as the business does.

THE NUMBERS BEHIND IT

12%

Median rise in cost per delivery — the second consecutive year at this level.

88%

Say delivery cost is growing as fast as revenue, or faster. Only one in eight is creating operating leverage.

55.7%

Name predictability or first-attempt success as the promise that matters most. Fastest possible delivery: 11.4%.

THE CENTRAL QUESTION OF THIS REPORT

How do you deliver the experience customers expect without letting the economics of fulfillment break?

INSIDE THE REPORT

Contents

01 **Growth has a delivery problem**

Growth is losing its delivery operating leverage.

02 **Certainty is the new premium**

Better delivery no longer means faster delivery.

03 **The false choice between CX and cost**

Poor CX and poor economics are symptoms of the same problem.

04 **Complexity is the hidden tax on growth**

Flexibility without control is just fragmentation.

05 **AI adoption is accelerating. AI advantage is not.**

Adoption is outpacing the ability to operationalize it.

06 **The delivery leaders playbook**

Five operating principles that follow from the findings.

07 **From cost center to commercial capability**

The closing argument.

Last Mile LEADERS

Hosted By FarEye

2026

FOREWARD

A NOTE FROM THE CEO'S DESK.

Cost, speed, and customer experience used to be a choice. In 2026, these are one concurrent expectation.

We surveyed 500+ Last-Mile Leaders and operators from enterprises across the region. They told us the classic trade-off between cost, speed, and customer experience has collapsed. The majority of respondents are now running multi-priority strategies, demanding all three from a single operating model.

The supply chain models and delivery promises made to customers have become dynamic, however, the systems behind these models and promises have not. This gap is the defining tension of 2026.

We hope it gives you a sharper read on where the market is going, and a clearer view of the work in front of all of us in the next twelve months.



KUSHAL NAHATA

CO-FOUNDER & CEO

FarEye

CHAPTER ONE

01

Growth has a delivery problem.

Growth is losing its delivery operating leverage. Costs are rising exactly when fulfillment is supposed to be helping the business scale profitably.

THE QUESTION WAS

How do we reduce logistics cost?

THE QUESTION IS NOW

How do we stop delivery cost from eating our growth?

1.1

Two years at 12%. This is the baseline now.

The typical delivery cost rose 12% in 2026, matching the increase reported in 2025. Costs are not settling back down.

Two consecutive years at this level tell you something a single year cannot: this is the new operating baseline, not a temporary spike to be absorbed and forgotten.

FIG. 1 — DELIVERY COST INFLATION, YEAR ON YEAR

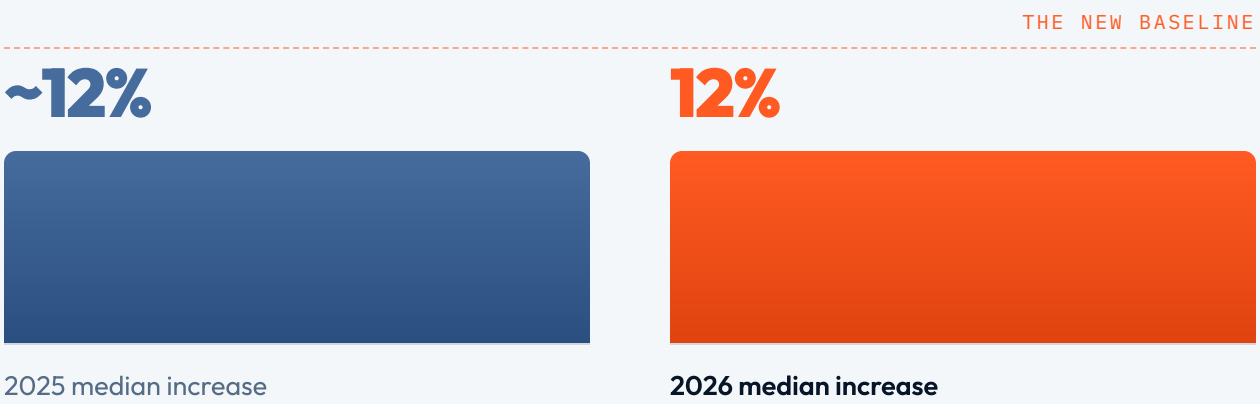

FIG. 2 — HOW THE 2026 INCREASE IS DISTRIBUTED



Six in ten operators report increases above 10%. One in five reports an increase above 20%.

1.2

Growth has stopped paying for itself.

Operating leverage is the simple test of whether scale is working: as volume grows, does the cost of serving each order fall? For most operators, the answer is now no.

FIG. 3 — DELIVERY COST AGAINST REVENUE GROWTH

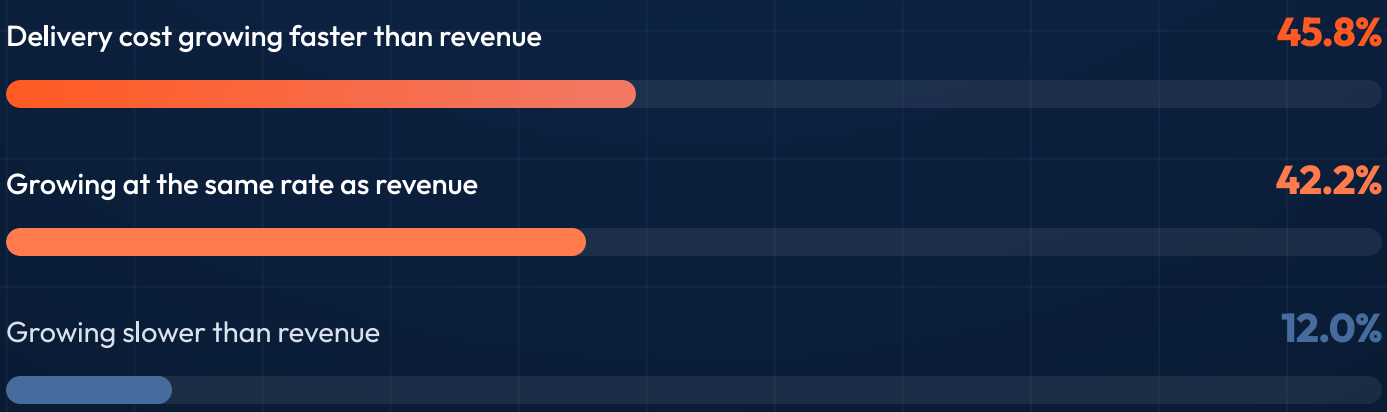

88% of operators are not generating delivery operating leverage as they grow.

An enterprise can grow its top line while quietly weakening the economics under every additional order.

1.3

The pressure is external. The opportunity is internal.

FIG. 4 — COST PRESSURES PLACED IN THE TOP THREE

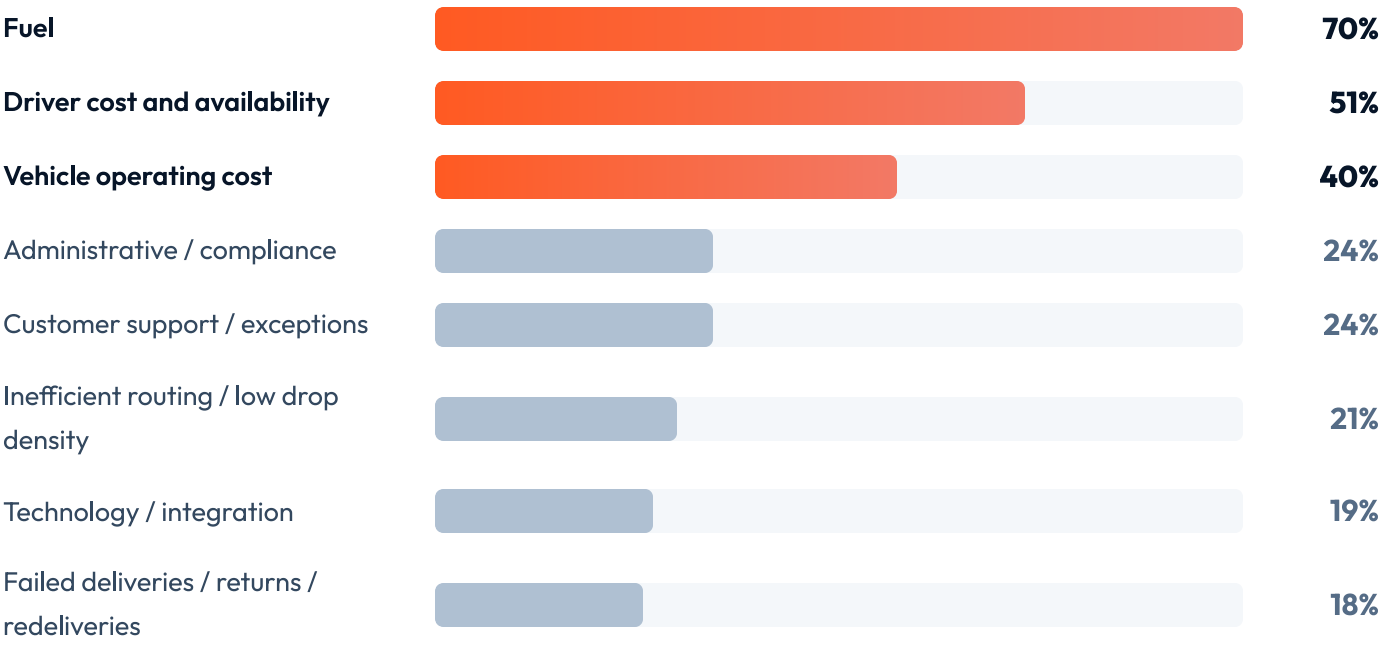

The three most-cited pressures are all external market inputs.

Fuel prices and labor markets cannot be engineered away. But the volume of fuel, labor, vehicle capacity and manual intervention consumed per delivery is entirely an internal decision — and that is where the recoverable margin sits.

> **The controllable response to uncontrollable inflation is operational productivity.**

ACTIONABLE TAKEAWAY

What this asks of the leadership team

If delivery cost is compounding against revenue, it has outgrown the function that owns it. Three things follow.

01

Make it a board metric.

Delivery economics belong in the same review as gross margin, not in a quarterly logistics update. Cost per delivery, on-time performance and exception volume should be reported to the board on the same cadence as revenue.

02

Get to cost-to-serve by segment.

Most operators can state total logistics spend. Far fewer can state cost to serve by promise tier, by geography, or by capacity pool — which is the level at which decisions actually get made.

03

Attack the exception, not the rate card.

Exceptions are where cost and customer experience fail together. Automated detection and resolution of at-risk deliveries is the fastest-payback application of AI in the last mile, and the one most likely to survive a CFO review.

CHAPTER TAKEAWAY

The last-mile challenge is no longer reducing today's cost. It is preventing delivery cost from scaling one-for-one with tomorrow's growth.

02

Certainty is the new premium.

Better delivery no longer means faster delivery. Last-mile competition has been framed around speed for a decade. Customers have moved on.

55.7% name predictability or first-attempt success as the promise that matters most. Fastest possible delivery: 11.4%.

2.1

Certainty beats speed, almost five to one.

FIG. 5 — THE DELIVERY PROMISE THAT MATTERS MOST TO CUSTOMERS

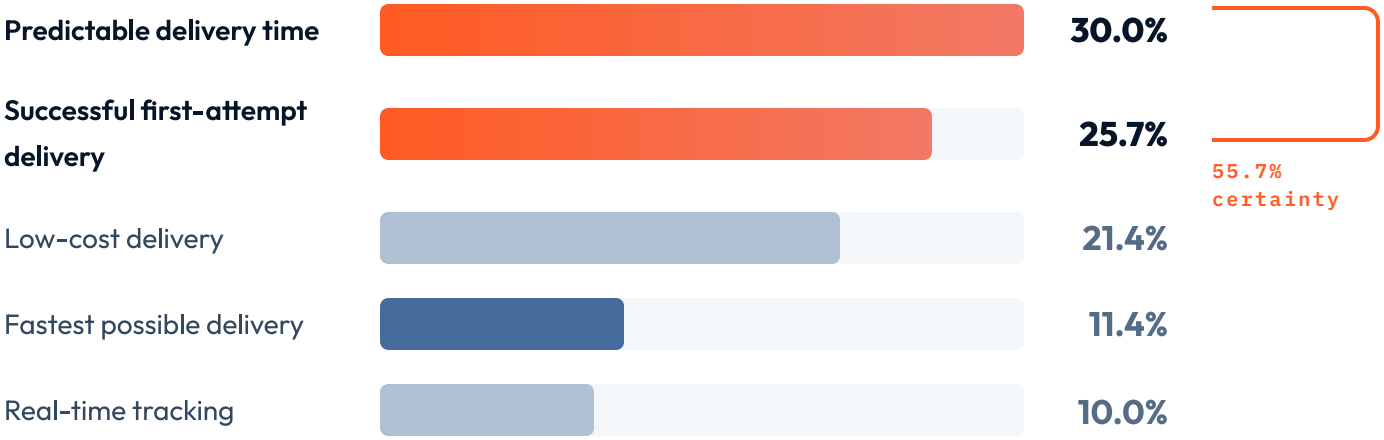

Predictability and first-attempt success together account for 55.7% of the promise operators say matters most. Fastest possible delivery: 11.4%.

The battleground has moved from velocity to certainty, and the distinction is economic. Speed is bought with excess capacity, expedited fulfillment and tighter operating windows — you pay for it every single day. Certainty asks a cheaper question: can the enterprise consistently do what it already said it would do?

2.2

Enterprises are promising certainty they cannot yet deliver.

FIG. 6 — THE PROMISE-CAPABILITY GAP, ACROSS THE SAMPLE

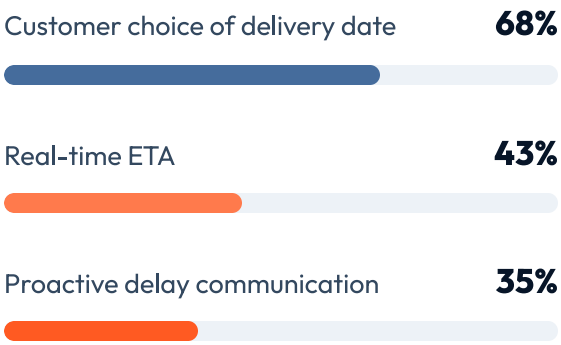

FIG. 7 — AMONG THOSE WHO SAY PREDICTABILITY MATTERS MOST

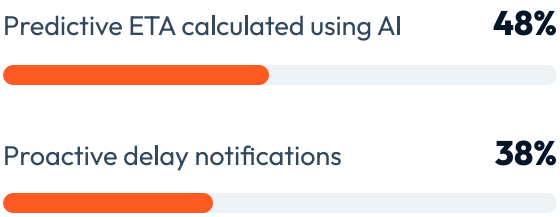

Base: operators who named predictable delivery time as the promise that matters most.

> **Fewer than half of the operators who say predictability matters most have the predictive ETA capability that makes a predictable promise possible.**

The industry has correctly read what the customer values. Execution has not caught up. A delivery date the customer chooses is a promise; without a predictive ETA underneath it, it is a guess with a calendar picker attached.

2.3

Companies digitized the happy path and left the failure state alone.

Look at which capabilities are live versus which are still on roadmaps, and a clear pattern appears. The tools that show a customer everything is fine are in production. The tools that let a customer fix something going wrong are next year's project.

FIG. 8 — CUSTOMER-FACING CAPABILITIES: IN USE AGAINST PLANNED

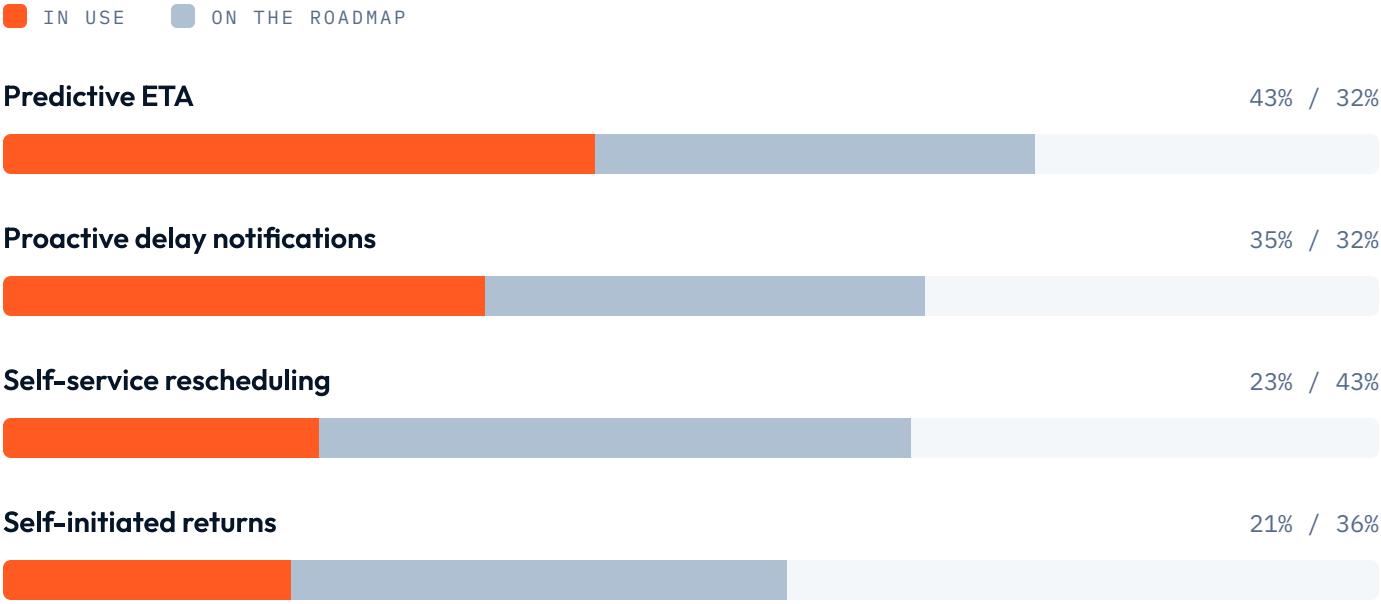

Customers can track an order. They cannot resolve a delivery that is going wrong. That gap gets expensive at precisely the moment the promise fails — when a self-service reschedule costing near-zero is replaced by a support call, a failed attempt, and a redelivery.

ACTIONABLE TAKEAWAY

The predictability ladder: where does your network sit?

The first benchmark worth setting is the simplest one, and most operators have never set it: what counts as “promise kept”? A dynamic two-hour window and a multi-day carrier SLA are both promises — they are not the same product, and they do not cost the same to keep.

Place your network against four observable stages.

FIG. 9 — THE PREDICTABILITY MATURITY LADDER (1 OF 2)

01 REACTIVE

Delivery window	Anytime within 2–5 days, on carrier contract SLAs
ETA accuracy	Opaque or unavailable for over 80% of orders; measured only on escalation
WISMO rate	Above 35% of orders
Customer engagement	Reactive: the customer calls support to find the order
Promise architecture	One promise for all customers

02 DEFINED

Delivery window	Day of delivery, or “today / tomorrow”
ETA accuracy	Visibility on 80–90% of shipments, measured periodically
WISMO rate	25–35%
Customer engagement	Static: status on demand or at milestones
Promise architecture	Two tiers (standard / express)

FIG. 9 — THE PREDICTABILITY MATURITY LADDER (2 OF 2)

03 PREDICTABLE

Delivery window	Shift-level window of 4–8 hours, limited view of sequence
ETA accuracy	Visibility on 90–95%, measured and monitored daily
WISMO rate	15–25%
Customer engagement	Predictive: the customer is told before the delay reaches them
Promise architecture	Three tiers, with cost-to-serve roughly known by tier

04 ARCHITECTED

Delivery window	Dynamic 2-hour windows, traffic-aware, varied by tier, zone and customer
ETA accuracy	Visibility on 95%+, measured per route and per carrier in near real time
WISMO rate	Under 15%
Customer engagement	Two-way: the customer can reschedule, redirect or release the ETA mid-route
Promise architecture	A portfolio of promises mapped to geography, capacity and margin

Position on predictability — but define the window, and invest in the visibility layer that makes the promise keepable before you advertise it.

DESIGN PRINCIPLE

Stop asking only “how fast should delivery be?” Start asking “which promise, to which customer, in which geography, through which capacity pool, at what margin?”

That question turns the last mile from a cost center into a managed promise architecture.

CHAPTER TAKEAWAY

The premium delivery experience is less about arriving fastest and more about removing uncertainty.

03

The false choice between CX and cost.

Poor customer experience and poor economics are not two problems. They are the same operational failure, measured at two different desks.

29.5
points of on-time performance separate the low and high control cohorts

2.4×
the cost inflation in the highest WISMO cohort against the lowest

7.1
points better on-time performance where delay notifications are proactive

The instinct in most organizations is to treat customer experience as something you buy with margin — better service costs more, cheaper service disappoints. The data points the other way. The operators with the worst customer experience also carry the worst cost trajectory, because both are produced by the same thing: deliveries that do not go as planned, and the manual work required to rescue them.

The goal is not maximum service at any cost, or minimum cost at the customer's expense. The strongest outcomes appear where enterprises reliably execute an appropriate promise.

3.1

Control, not investment, separates the outcomes.

Control here means something specific: the ability to see, decide and act across the whole network in near real time — regardless of who owns the vehicle. Not ownership of assets, and not ownership of software. Command of the operation.

Split the sample by reported control and the results diverge sharply, at identical levels of investment.

CONTROL LEVEL

Low

CONTROL LEVEL

Medium

CONTROL LEVEL

High

FIG. 10 — OUTCOMES BY CONTROL COHORT

Reliability rises, customer contact falls, and cost inflation drops — all at the same time, for the same money.



Based on the 41 respondents who answered this question. Median technology investment is identical across all three cohorts.

Read the three columns together. Reliability rises, customer contact falls, and cost inflation drops — all at the same time, for the same money. Nearly thirty points of on-time performance separate the low and high cohorts. Investment does not explain operational control. Neither does technology maturity.

The best delivery economics begin with keeping the promise, not with cutting the service.

3.2

Control changes the problems you get to have.

Operators with higher control do not just perform better. They report a materially different set of problems with the same outsourced partners — and the three that fall hardest are precisely the CX complaint, the data complaint and the cost complaint.

FIG. 11 — OUTSOURCED-FLEET CHALLENGES, BY CONTROL LEVEL

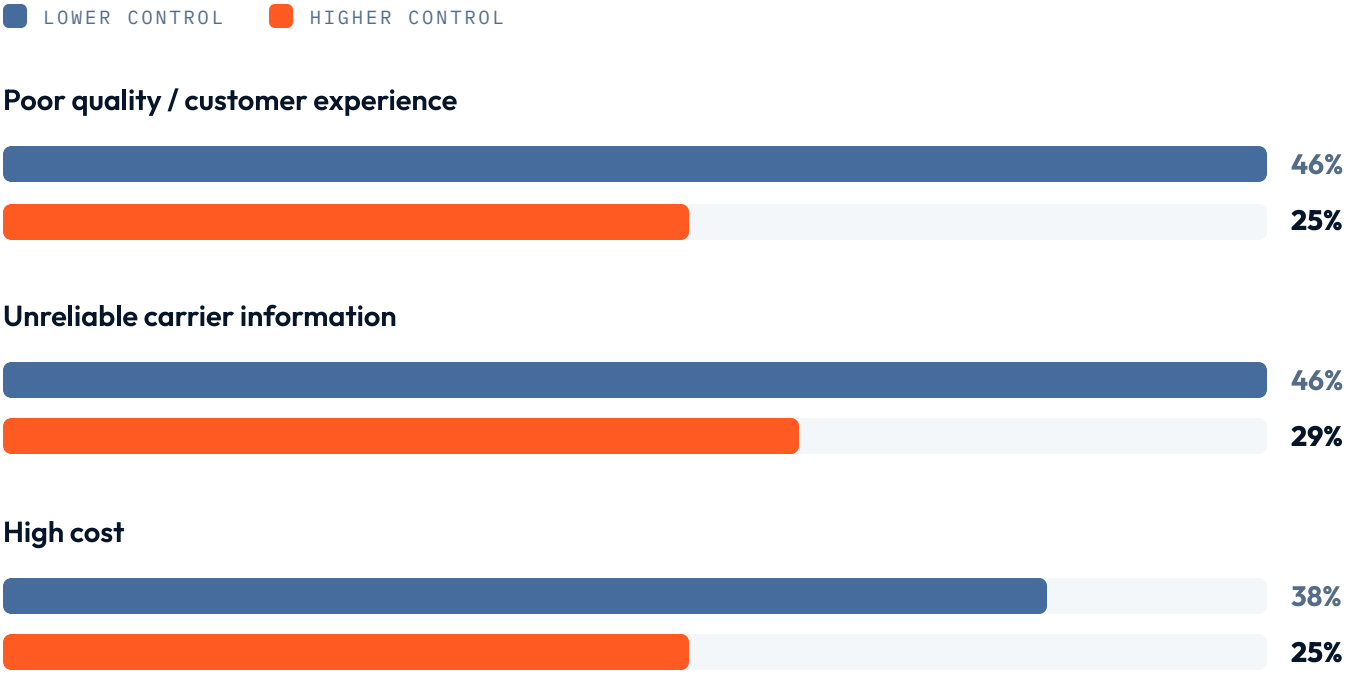

THE QUESTION MOVES FROM

What is happening out there?

TOWARD

How do we manage this partner against the outcome we agreed?

That is the shift worth engineering for. Low-control operators spend their week discovering what went wrong. High-control operators spend it managing partners against an agreed outcome — which is a cheaper week and a better customer experience at the same time.

ACTIONABLE TAKEAWAY

WISMO behaves like an economic indicator

“Where is my order?” contacts are usually owned by customer service and reported as a contact-center statistic. Read across the cohorts and it behaves like something else entirely — a live index of delivery friction.

FIG. 12 — OUTCOMES BY WISMO COHORT

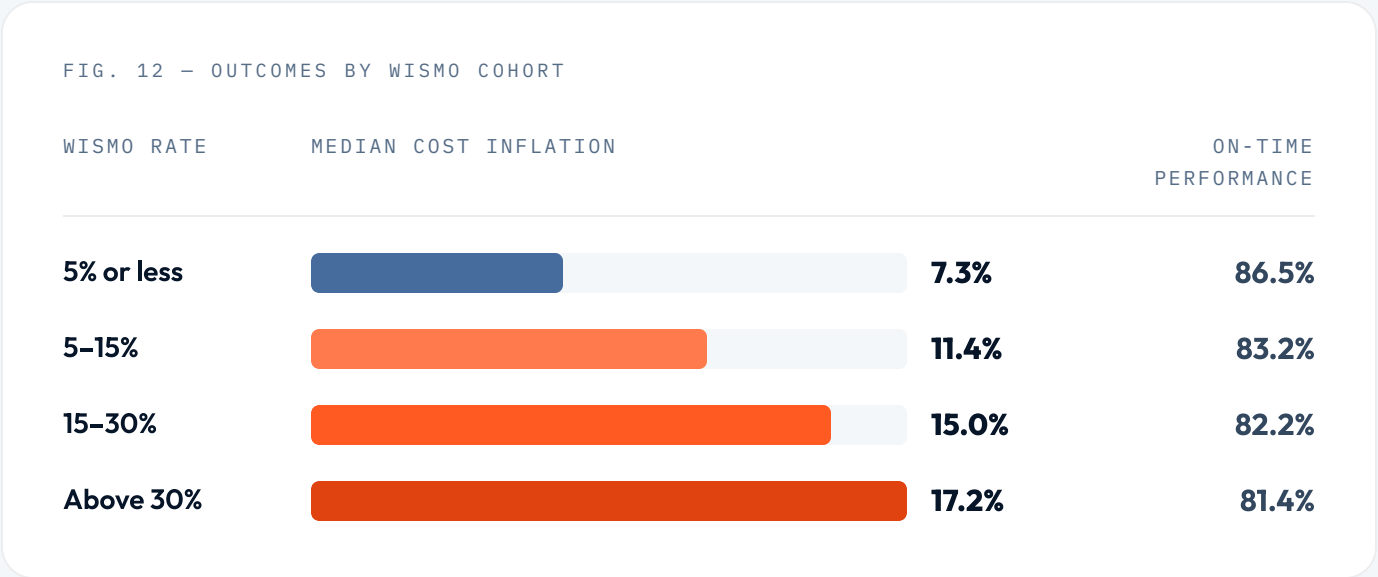

Cost inflation more than doubles between the lowest and highest WISMO cohorts — 7.3% against 17.2% — while on-time performance falls in step. Every WISMO contact is a customer who was not confident, a delivery that was not visible, and a person paid to close a gap the operation created.

WHAT TO DO WITH THIS

Instrument WISMO as a leading indicator and review it monthly alongside cost per delivery. When the WISMO rate moves, it is telling you where your cost per delivery and your promise reliability are heading next — usually before either shows up in the P&L.

ACTIONABLE TAKEAWAY

The cheapest CX intervention is also an operational one

Proactive delay notification is the only customer-experience capability in this study associated with better on-time performance — not just with happier customers.

FIG. 13 — ON-TIME PERFORMANCE AND PROACTIVE DELAY NOTIFICATIONS

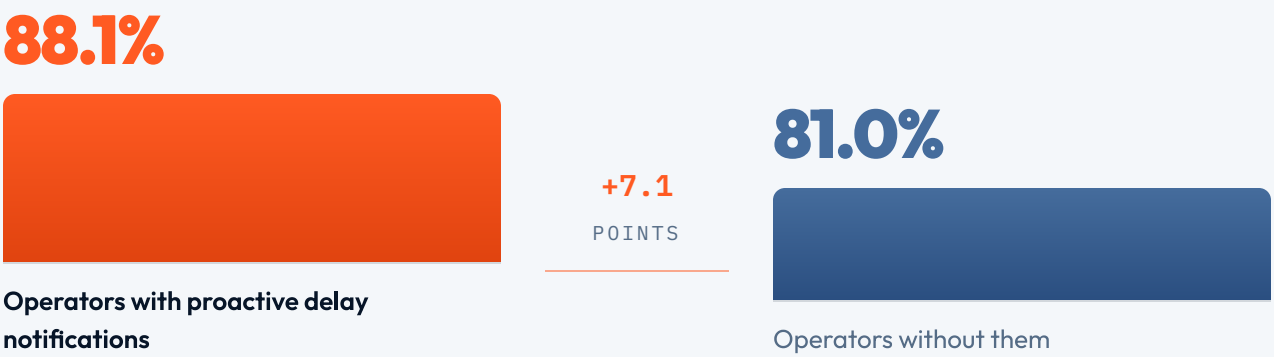

A seven-point gap in on-time performance, from a capability usually budgeted as communications. The mechanism is straightforward: telling a customer early converts a failed attempt into a rescheduled one, and a rescheduled delivery is a delivery that still counts. Communication stops being a customer-service layer and becomes part of operating performance.

> **Better customer experience does not require a more expensive promise. Often it just requires managing uncertainty better.**

CHAPTER TAKEAWAY

Better CX and better economics reinforce each other when both are built around reliable execution. Fix the failure state and you fix both ledgers at once.

NEXT

Maintaining reliability gets harder as the network expands. The next pressure on profitable delivery is not scale. It is fragmentation.

04

Complexity is the hidden tax on growth.

More volume is not the problem. Growth introduces more carriers, more fleets, more systems, more markets and more operating rules — and every new source of flexibility creates another point of fragmentation.

What you introduce to create flexibility usually ends up creating fragmentation.

4.1

The hybrid network is already the default.

FIG. 14 — FLEET COMPOSITION AND OUTSOURCING INTENT

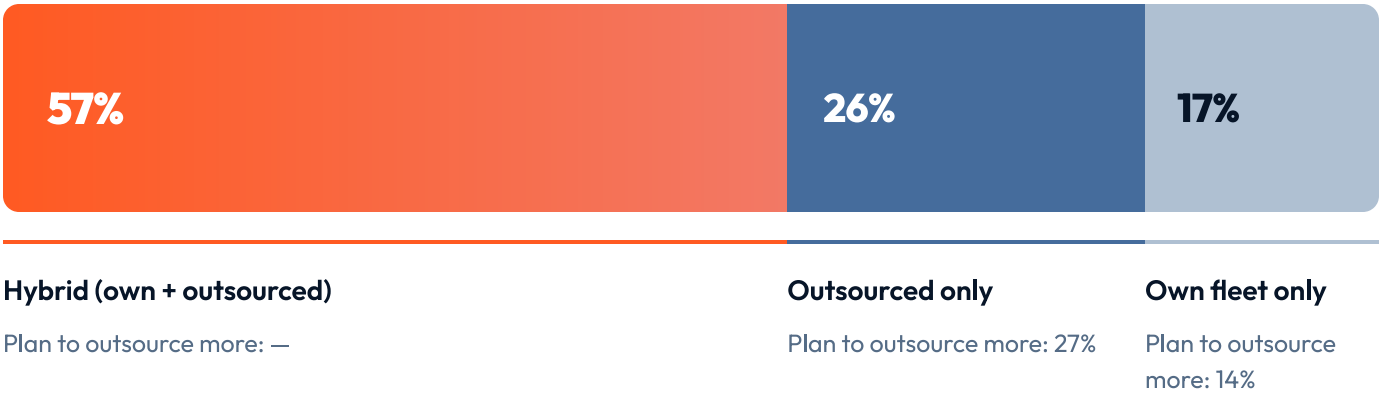

57% already run a hybrid network — and more than half of the remainder plan to outsource more.

The dominant model is neither pure ownership nor pure outsourcing. Enterprises are building mixed networks in pursuit of flexibility, which raises the stakes on a hard problem: keeping one promise consistent across parties the enterprise does not own.

4.2

The Goldilocks fleet: not too owned, not too outsourced.

The old fleet debate was simple — own it for control, outsource it for flexibility. The data suggests that debate is losing relevance. Hybrid operators report roughly 11% higher on-time delivery and close to half the WISMO rate of single-model networks. And across every model type, close to half of operators plan to deepen carrier relationships rather than replace them.

Both extremes are unstable, for different reasons, and 2026 operating conditions punish both:

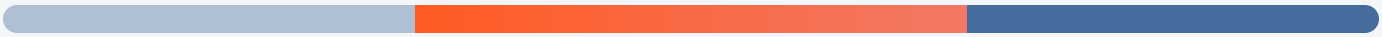

PURE OWN FLEET | THE GOLDILOCKS ZONE | PURE OUTSOURCED

Pure own-fleet does not survive demand variance.

Volume volatility is now structural. Capacity sized for peak is idle capital for ten months of the year.

Pure outsourced does not survive experience pressure.

When the promise is certainty and the data belongs to someone else, you cannot keep what you sold.

Between them sits the Goldilocks Zone: enough owned or dedicated capacity to guarantee the critical SLAs, enough flex capacity to absorb the variance. Not too owned. Not too outsourced. Just orchestrated enough.

ACTIONABLE TAKEAWAY

Hybrid in name, homogeneous in behavior

The test is not what your network looks like on a slide. It is whether — at any moment, for any parcel, in any capacity type — your team can answer four questions without picking up a phone.

44% Where is it?

44% of operators cannot answer this today. Lack of real-time visibility is the single most-cited operating bottleneck. The minimum bar is milestone, GPS, ETA and exception alerts across every capacity type.

35% Why this carrier?

35% cite carrier performance as a technology gap, because allocation today runs on static rules and habit. Scorecards have to feed allocation decisions, not just quarterly business reviews.

34% What is it costing me?

34% cite audit and reconciliation as a bottleneck. Margin leaks quietly through manual invoice review. Automated audit rules and exception-based review are now table stakes.

18% What will it deliver?

Only 18% flag dynamic allocation as a challenge — which is itself the signal. Most operators are not yet mature enough to know what static carrier rules are costing them.

> **An operator who cannot answer these is not running a hybrid network. They are running several networks and hoping.**

CHAPTER TAKEAWAY

Flexibility without a control layer is just fragmentation with better branding. Standardize the layer that governs outcomes — not the way every carrier operates.

05

AI adoption is accelerating. AI advantage is not.

2026 is the year AI moved from experimentation into implementation. Adoption is no longer the useful measure — the divide is between enterprises that own AI capability and those that can operationalize and scale it.

THE QUESTION WAS

Will delivery organizations adopt AI?

THE QUESTION IS NOW

What happens after adoption?

5.1

Adoption moved twenty points in twelve months.

FIG. 15 — AI ADOPTION, 2025 TO 2026

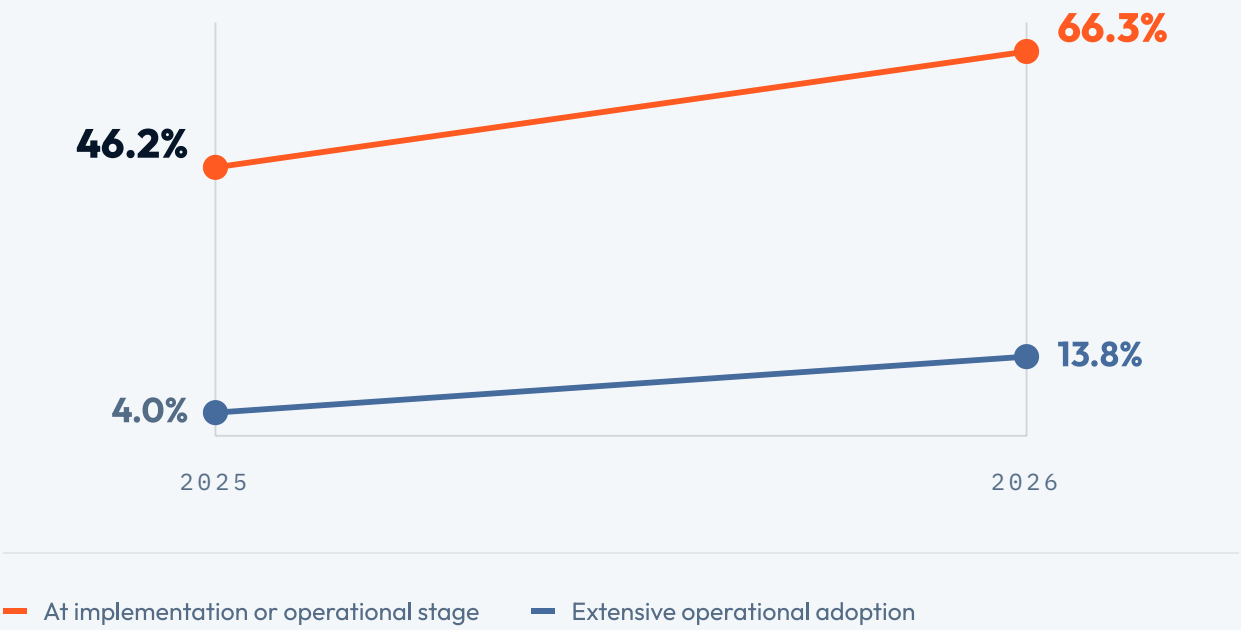

A twenty-point move in a single year, and close to a fourfold increase in extensive operational adoption. The market has moved decisively past awareness. What it has not resolved is what to do with the capability once it is switched on.

5.2

AI is trusted to advise, not yet to act.

FIG. 16 — LEADING AI USE CASES

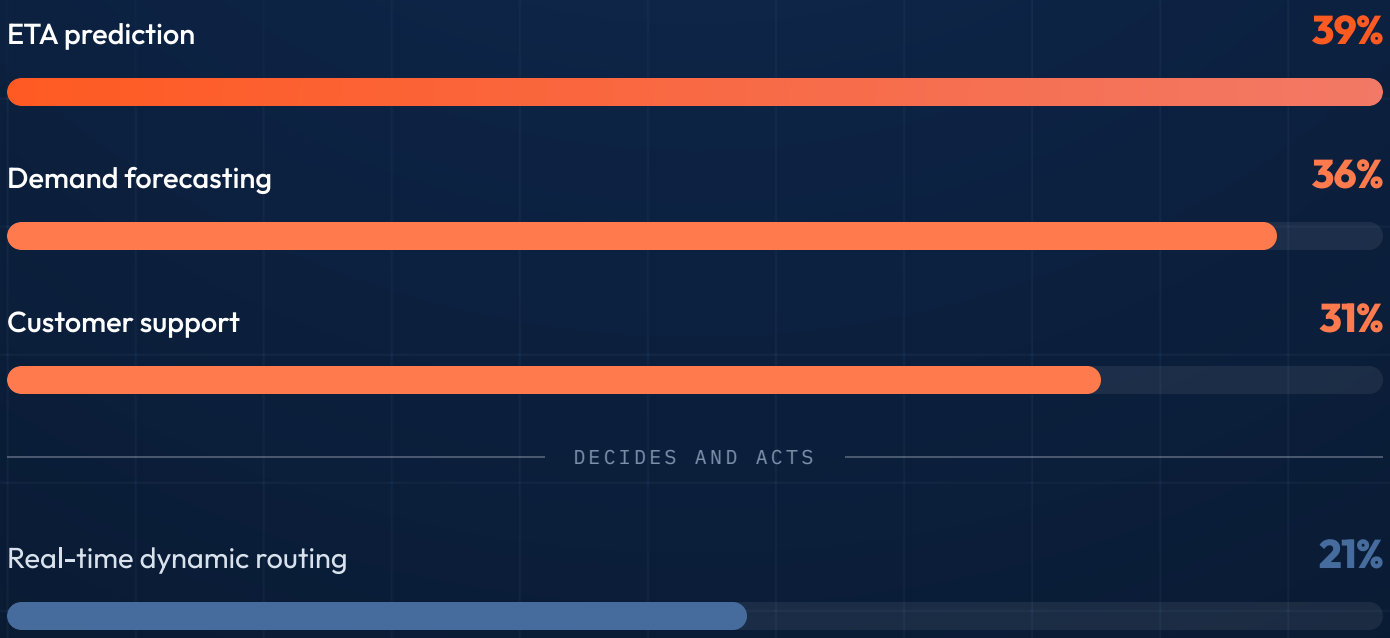

The top three use cases predict or explain. The one that decides and acts — real-time dynamic routing — sits well behind. Enterprises are comfortable asking AI what will happen. They are considerably less comfortable letting AI decide what to do next.

The gap between 39% and 21% is not a technology gap. It is a trust gap — and trust is an engineering problem with a governance answer.

ACTIONABLE TAKEAWAY

What turns adoption into advantage

Across the operators seeing measurable returns, three things are consistently in place before the model is — and one thing is deliberately designed after it.

01 **A data layer that can be trusted.**

Models inherit the quality of the operation underneath them. Standardized milestones across carriers, a common exception taxonomy and reconciled cost data are the unglamorous prerequisites. An ETA model trained on inconsistent carrier scans will produce confident, wrong answers faster than a human could.

02 **A decision it actually feeds.**

A forecast that lands in a dashboard changes nothing. Value appears when the output triggers an action inside the workflow — a reallocation, a customer notification, a re-sequenced route — without a person copying it into another system.

03 **Human-in-the-loop by design.**

Human-in-the-loop is not a safety blanket to be removed later. It is the mechanism that generates the evidence for widening autonomy: which decisions the model gets right, where it drifts, and which exceptions must always escalate. Start with low-risk, high-frequency decisions and let the override log tell you where the ceiling is.

04 **Governance built in parallel, not retrofitted.**

Explainability, autonomy thresholds, a documented override path, audit trails and ROI gates. Every enterprise will need this eventually — the ones that build it alongside the intelligence rather than after an incident will scale autonomy years earlier than the ones that don't.

ACTIONABLE TAKEAWAY

The five stages of AI maturity

Most operators are somewhere in stages two and three. The move that matters is not the next model — it is the next stage.

FIG. 17 — THE AI MATURITY LADDER (1 OF 2)

Stage 1 REACTIVE

WHAT IT LOOKS LIKE

Manual planning, fragmented tools, human firefighting. The signature symptoms are an unknown cost per delivery, high WISMO and inconsistent carrier data.

HOW TO MOVE UP

Focus on measurement, not technology. Establish cost, service and exception baselines before you buy anything.

Stage 2 DIGITIZED

WHAT IT LOOKS LIKE

Tracking and dashboards exist; decisions have not caught up. Visibility is necessary but not sufficient.

HOW TO MOVE UP

Fix the data before the intelligence. Standardize milestones across carriers, adopt one exception taxonomy, and reconcile cost data to a single source so the numbers on the dashboard can be acted on.

Stage 3 OPTIMIZED

WHAT IT LOOKS LIKE

Routing, dispatch and SLA management all work — function by function, each tuned to its own KPI. The seams show at the joins.

HOW TO MOVE UP

Bring promise design, routing, allocation and customer communications into a single operating layer that shares context end to end.

FIG. 17 — THE AI MATURITY LADDER (2 OF 2)

Stage 4 ORCHESTRATED

WHAT IT LOOKS LIKE

A unified operating layer surfaces decisions, not just data. Exceptions get pre-empted. Capacity reallocates dynamically across own, outsourced and gig pools.

HOW TO MOVE UP

Identify low-risk, high-frequency decisions and automate them first, with the override path instrumented from day one.

Stage 5 ADAPTIVE

WHAT IT LOOKS LIKE

Real-time orchestration with selective autonomous execution. Capacity, promises and exceptions adapt continuously to real demand and real performance, built on human-in-the-loop.

HOW TO MOVE UP

The work shifts from running the operation to governing it: autonomy thresholds, model drift, ROI gates and the human override path.

STAGE 1 — STAGE 5

CHAPTER TAKEAWAY

Adoption is now table stakes. The advantage belongs to enterprises that can move AI from forecasting the operation to governing it — safely, and with evidence.

CHAPTER SIX

06

The delivery leaders playbook.

Five operating principles for turning delivery into a commercial capability. Each follows directly from the findings.

01 Design the promise around what customers value.

02 Treat reliability as both a CX metric and a cost metric.

03 Build control without giving up flexibility.

04 Measure technology by the decisions it improves.

05 Move AI from forecasting the operation toward governing it.

PRINCIPLE 01

Design the promise around what customers value, not what competitors advertise.

WHY

Maximum service is not better service. Leaders decide which promises matter enough to justify the cost of keeping them.

WHAT TO DO

Manage the delivery promise as an economic decision.

```
Service value × probability of successful execution × cost-to-serve.
```

PRINCIPLE 02

Treat reliability as both a CX metric and a cost metric.

WHY

Customer experience and cost-to-serve share the same operational root cause.

WHAT TO DO

Bring both into one performance conversation, reviewed by one owner.

```
OTIF + first-attempt rate + WISMO + exceptions + cost per delivery.
```

PRINCIPLE 03

Build control without giving up flexibility.

WHY

Delivery networks will keep becoming more distributed. Leaders need central control without forcing every partner into a single operating model.

WHAT TO DO

Standardize the layer that governs outcomes, not the way every carrier operates.

> Performance rules, data standards, decision logic, customer promises.

PRINCIPLE 04

Measure technology by the decisions it improves, not the capabilities purchased.

WHY

The unit of technology value is no longer feature adoption. It is improved operational execution.

WHAT TO DO

Evaluate every platform against the operating outcomes it changes.

> Decision latency, exception reduction, reliability gain, time to deploy, cost-to-serve.

PRINCIPLE 05

Move AI from forecasting the operation toward governing it.

WHY

The next stage of AI maturity is not another prediction model. It is the controlled automation of operational decisions.

WHAT TO DO

Build the governance layer alongside the intelligence, not after it.

```
Explainability + autonomy thresholds + human override + auditability + exception escalation.
```

THE PLAYBOOK IN ONE LINE

Decide which promises are worth keeping, build the control to keep them, and measure every investment by the decisions it improves.

07

From cost center to commercial capability.

Delivery belongs in the commercial conversation.

For most enterprises, delivery cost is no longer scaling favorably against revenue. At the same time, the experience customers value is moving away from maximum speed toward reliability. And the failures that weaken that experience carry higher operational friction and worse cost trajectories with them.

Those three findings are usually treated as three separate agendas — a finance problem, a marketing problem and an operations problem. They are one problem, and it has one answer.

The opportunity is not to choose between customer experience and cost efficiency. It is to redesign the operating model so that the same capabilities improve both: a deliberate promise, executed reliably, across a network the enterprise can actually command.

THE NEW LAST-MILE ADVANTAGE

The winners will not be the companies that promise the most. They will be the companies that keep the promises that matter, at economics the business can scale.

About this research.

Eye on the Last Mile America 2026 surveyed delivery and supply chain leaders at enterprises across the United States, capturing more than 3,000 data points on delivery economics, the customer promise, network design and technology adoption. Fieldwork ran across Q1 and Q2 2026.

Where a finding is based on a subset of respondents rather than the full sample, the base is stated alongside the figure.



3,000+
data points on delivery economics, promise, network and technology

Q1–Q2
2026 fieldwork window

USA
delivery and supply chain leaders at enterprises

PUBLISHED BY

Eye on the Last Mile is published by FarEye, host of Last Mile Leaders America 2026 — three days of closed-door strategy for the executives who own delivery.

Last Mile LEADERS

Hosted By FarEye

LAST MILE LEADERS · HOSTED BY FAREYE

26–28 AUGUST 2026 · THE GWEN, CHICAGO, USA

Closed-door strategy for the executives who own delivery.