Deploying Autonomous Trucks in Middle-Mile Logistics: A Case Study of a 3PL and AV Startup Partnership

Ashvin Kulkarni
Senior Technical Product Manager, Collaborative Innovations

DOI: 10.29322/IJSRP.14.04.2024.p14XXX

Abstract- Autonomous vehicles have been talked about for years, but actual commercial deployments at scale in logistics remain rare. This paper documents one such deployment: a major 3PL logistics company moving from a five-truck pilot to a 54-truck autonomous fleet through its partnership with an AV startup specializing in middle-mile logistics. We describe the financial logic, contract structure, and real operational outcomes that drove the decision. The numbers are straightforward: the AV startup's trucks cost $32.30 per stop against the incumbent courier's $43.82, and over a 60-month contract the projected savings reach $6.5 million. Less obvious—and more interesting—are the contract provisions that make this work without betting the farm on unproven technology, and the early entry delivery improvements that turned out to matter more to clients than anyone initially expected.

Index Terms- autonomous vehicles, middle-mile logistics, D2DDU delivery, cost-per-stop, early entry percentage, fleet economics, supply chain sustainability

I. Introduction

The freight industry has a middle-mile problem. Long-haul trucking has garnered most of the autonomous vehicle attention, and last-mile delivery gets the consumer headlines—but the segment connecting distribution hubs to regional post offices, retail backrooms, and fulfillment centers sits in a quieter and arguably more tractable space for automation. Routes are fixed. Stops are predictable. Schedules repeat. These are not the chaotic conditions that make autonomous driving hard.

The AV startup at the center of this study identified this gap and built specifically for it. Their Class 3–6 box trucks are not trying to navigate a chaotic urban intersection or handle an unexpected merge on the interstate at 70 mph. They run the same routes, at consistent speeds, to the same loading docks, every day. That repetition is actually an asset for machine learning—the more times a system sees the same route, the more confident its decision-making becomes.

The 3PL logistics company began a modest pilot of this technology in the Dallas-Fort Worth market with five trucks serving 74 USPS Destination Delivery Units (DDUs). What they found changed the investment calculus quickly. Not only were costs lower, but on-time early entry rates climbed from 80% to 98%. That kind of improvement—getting more parcels into USPS hands before the 8 a.m. cutoff—directly translates to faster consumer delivery. Clients noticed.

This paper traces the decision from pilot to a 54-truck expansion across three U.S. markets, examining the financial modeling, contract architecture, service-level evidence, and environmental case that together supported a $26 million minimum commitment. Our aim is to give logistics operators a realistic picture of what AV adoption in middle-mile actually looks like right now—not the theoretical version, but the messy, contract-negotiated, risk-mitigated reality.

II. The AV Startup and the Middle-Mile Market

The AV startup operates in a segment with no direct autonomous competitor. That sounds like a bold claim, and it is worth being precise about why it holds: while other companies have pursued long-haul freight or broad-market AV platforms, none is focused on the short-haul, fixed-route, B2B middle-mile niche this startup occupies. Its existing customer base at the time of the 3PL's expansion decision included several major U.S. and Canadian retailers and consumer goods companies, representing a substantial and growing enterprise-grade commercial footprint.

The truck itself is a Class 6 medium-duty 26-foot box vehicle, diesel-powered, with a daily range of 300 miles over a 12-hour operation window. It carries LiDAR, cameras, and sensor arrays sufficient for Level 4 autonomous operation on pre-mapped fixed routes. Each vehicle includes an Autonomous Vehicle Operator (AVO)—a safety driver provided and paid by the AV startup—who remains on board until regulatory and operational conditions support full driver-out operation. Critically for the customer: the AVO is the AV startup's cost, not the 3PL's.

Manufacturing lead time is 12 months. That constraint shaped the urgency of the 3PL's investment decision more than any financial model: with the AV startup's order book filling and no competitor capable of similar service, waiting another year meant either accepting a spot further back in the queue or paying higher prices once the current contract window closed.

III. Pilot Evidence: What the DFW Deployment Actually Showed

Before committing to 49 additional trucks, the 3PL had roughly 18 months of live route data from its DFW pilot to lean on. Five trucks. Seventy-four DDUs. Real routes, real weather, real variability in USPS dock conditions. The headline figures held up.

Cost per stop: the AV startup ran at $32.30 versus the third-party courier at $43.82. That is a 26% reduction on a metric that compounds across hundreds of thousands of annual stops.

Early Entry Percentage: 98% of DDU deliveries arrived before the 8 a.m. USPS cutoff under the AV startup, compared to 80% for the courier. The 18-point gap is not trivial. Every parcel that crosses the threshold before 8 a.m. is eligible for same-day carrier sortation and next-day consumer delivery. Miss that window and the parcel sits, adding a full day to the cycle. The 3PL's clients ship at high enough volumes that this difference showed up in their delivery metrics.

Scanning accuracy: the pilot logged 99% accuracy on both delivery and return parcel scans. Incumbent couriers were hitting 95%. This matters for a different reason—scan data drives billing reconciliation, returns processing, and USPS service credit calculations. A 4-point improvement in accuracy reduces disputes and processing delays downstream.

No safety incidents occurred over more than one million miles driven across the AV startup's commercial fleet. For a company whose logistics operations touch the USPS—a government customer with zero tolerance for liability exposure—that record was a prerequisite for expansion, not just a nice-to-have.

IV. Financial Analysis

A. Cost Modeling Approach

The financial model compares two scenarios over a shared time horizon: the AV startup's lease arrangement versus a do-nothing baseline of continuing with third-party couriers. Courier costs were inflated at a conservative 2% year-over-year—conservative because actual transportation inflation averaged 4.5% annually between 2016 and 2022, including 14.6% in 2021 and 15.5% in 2022 [7]. If that kind of volatility recurs, the baseline cost trajectory shifts upward sharply and the AV startup's economics look even better.

The AV startup's pricing moves the other direction: contractual discounts of 2%, 3.5%, 5%, and 5% apply in years 2 through 5 respectively, so the gap between the two cost lines widens each year. A fuel surcharge of approximately 5.5% is layered onto the base lease rate, reflecting diesel price exposure, and is applied consistently to both scenarios since both involve diesel vehicles.

Excluded from the model: fleet tracking hardware (the AV startup uses its own telemetry), toll costs (current routes are toll-free), vehicle maintenance (entirely the AV startup's responsibility under the MSA), insurance and licensing (bundled in the per-truck fee), and incremental shipment costs from vehicle failure (the AV startup contractually provides replacement capacity the same day).

B. Summary Financials

Table I: 5-Year Cost Comparison – AV Startup vs. Courier Baseline (Values in $MM)

	2024 (6mo)	2025	2026	2027	2028	After	Total
Expense – AV Startup	(3.084)	(9.928)	(9.781)	(9.394)	(8.924)	(5.896)	(47.006)
Expense – Courier Baseline	(3.273)	(10.355)	(10.562)	(10.773)	(10.988)	(7.595)	(53.546)
Incremental Savings	0.189	0.427	0.781	1.379	2.065	1.699	**6.539**

Two commitment windows were modeled. The minimum—32 months including the 6-month notice requirement—yields $1.792MM in incremental EBIT and an NPV of $1.067MM. The full 60-month term produces $6.539MM in incremental EBIT and an NPV of $3.330MM. There is no capital investment in either scenario; the entire arrangement runs as an operating lease, with zero assets on the 3PL's balance sheet.

IRR and formal payback period were not computed. This is a cost-avoidance investment, not a capital project—the right frame is not 'when do we get our money back' but 'how much cheaper is this than what we're doing today.' The answer is: cheaper from day one, and increasingly so each year.

V. Deployment Across Three Markets

The 54-truck fleet rolls out across three U.S. markets: Dallas-Fort Worth (DFW), Ontario (ONT), and Stockton (SCK). DFW launches in July 2024 and includes the 5 existing pilot trucks; ONT and SCK both go live in October 2024.

Table II: Fleet Deployment by Market

Market	Trucks	Daily Stops	Routes	DDUs/Truck/Day	Launch
Dallas-Fort Worth	14	190	14	13.6	Jul 2024

Ontario (ONT)	25	291	25	11.6	Oct 2024
Stockton (SCK)	15	189	15	12.6	Oct 2024
Total	**54**	**670**	**54**	**12.4 avg**	—

Each route is served by one truck. Operations cover 100% of D2DDU routes in these three markets—this is not a hybrid model where the AV startup supplements couriers, it is a full replacement. That simplifies dispatch and eliminates the coordination overhead that comes with split coverage.

The 3PL's operations staff load the trucks exactly as they load any other vehicle today. No additional training is required. The AVO—the AV startup's responsibility—handles everything once the doors close.

VI. Service-Level Performance

Early Entry Percentage turned out to be the metric that resonated most with clients during the pilot. The logic is intuitive: a parcel that reaches a DDU before 8 a.m. can enter the USPS sortation process in time for same-day dispatch, giving it a realistic shot at next-day consumer delivery. Parcels arriving after 8 a.m. miss that window entirely.

Table III: Projected Early Entry Performance by Market

Market	Trucks	Stops	AV Startup EE%	Prior EE%	Delta
Dallas-Fort Worth	14	190	64%	51%	+26%
Ontario (ONT)	25	291	95%	95%	0%
Stockton (SCK)	15	189	83%	74%	+12%
Totals	**54**	**670**	**83%**	**77%**	**+8%**

Ontario shows no improvement because the incumbent courier was already hitting 95%—the market was already well-served, and the AV startup's benefit there is purely financial. DFW shows the largest gain (26 percentage points) because the prior baseline was poor, and the fixed-route consistency of autonomous operation closes the gap effectively. Stockton splits the difference.

Beyond EE%, the contracted SLA framework tightens on every dimension compared to standard courier agreements. Delivery and returns scanning accuracy moves from 95% to 99%. Truck availability is guaranteed at 98%. On-time delivery and returns targets hold at 95%. When SLAs are missed, the AV startup has 30 days to remediate outside peak periods (10 days during peak)—and if the gap is not closed, the contract is void.

One less-discussed benefit is data visibility. Currently, pallets are not scanned at USPS docks by courier drivers—the 3PL only sees return volume when parcels physically arrive at the hub and pass through the scan tunnel. The AV startup's fleet platform logs timestamped container-level scan data at every unloading event, shared in real time. That means the 3PL's operations team knows what is coming back, and when, before the truck arrives. Labor allocation becomes predictable. That is not a headline number, but it is real operational value.

VII. Contract Structure and Risk Mitigation

Autonomous vehicle contracts at this stage of market development are unusual documents. The technology is real and commercially proven, but regulatory environments are still settling, driver-out timelines are estimates rather than guarantees, and pricing dynamics could shift if the conventional courier market softens. The MSA between the 3PL and the AV startup was negotiated to address all three of those uncertainties directly.

A. Termination Provisions

There is no termination for convenience in the first 26 months. That is the commitment period—the 3PL accepted it in exchange for locked-in pricing and manufacturing priority. Starting at month 12, however, the 3PL gains a meaningful lever: it can present a bona fide third-party courier proposal, and the AV startup must either match the pricing or allow the 3PL to exit the affected routes. That provision essentially caps the contract's downside: if courier rates drop significantly and the AV startup won't follow, the 3PL can walk. From month 26 forward, either party can exit with 6 months' notice, no fees.

Volume decline is also addressed explicitly. A 60% drop in parcel volume on a given route, or a 50% drop in stop count, triggers a renegotiation process. If the parties cannot agree on reallocation or pricing adjustment, either party can terminate that route's AV. This provision protects the 3PL against committing capacity it no longer needs if client volume shifts.

Breach has a clear remedy path: written notice, a cure period (30 days outside peak, 10 days during peak), and contract termination if the issue is not resolved. The AV startup will provide additional trucks at its own cost to recover schedule during a performance improvement period—that clause shifts the remediation burden entirely to the vendor.

Transferability is permitted to affiliates and in connection with corporate transactions including mergers and acquisitions—without requiring the AV startup's consent. For a large publicly traded logistics company navigating an active M&A environment, this provision removes a common deal-complicating restriction.

B. Autonomous Technology and Regulatory Risk

Two scenarios were stress-tested. First: what if driver-out is delayed because the AV startup's software isn't ready, or because route complexity creates edge cases that take longer to resolve? The answer under the current MSA is that the AV startup absorbs the cost entirely. Pricing is not linked to driver-out status. The AVO stays on board, the AV startup pays for it, and the 3PL continues receiving the contracted price. The commercial arrangement was never contingent on immediate driverless operation.

Second: California's regulatory framework for Class 6 AV operation remains in progress. If legislation does not advance, electric trucks may be required under California Clean Fleet mandates (10% of fleet electric from 2025). Electric variants of the AV truck have a shorter range—200 miles versus 300—which could require route redesign to maintain coverage. The 3PL's response plan is to work with its Transportation Network Strategy team to optimize routes within the same truck count. No additional trucks would be needed, but the routing math changes.

VIII. Environmental Case

Autonomous trucks drive more efficiently than human drivers. The reasons are mechanical rather than philosophical: consistent speed with no unnecessary acceleration or braking, programmatic idle management at stops, and real-time tire pressure optimization. The AV startup's own fleet data shows 9% better fuel economy in AV mode versus manual operation on identical routes. The broader AV literature from the University of California San Diego and the National Renewable Energy Laboratory puts the industry range at 8%–10% for heavy-duty vehicles [5, 6].

Table IV: Annual CO2 Avoidance – 54 Trucks

Metric	Legacy	AV Fleet	Savings
CO2 per mile (kg)	1.18	1.08	0.10
Avg. miles per truck per year	65,602	65,602	—
CO2 per truck per year (kg)	77,599	70,995	6,604
Total CO2 per year – 54 trucks (kg)	**4,190,338**	**3,833,714**	**356,625**

Over 356,000 kilograms of CO2 avoided per year across the fleet. That figure is not the primary reason the 3PL made this investment, but it is a real number that feeds into corporate carbon neutrality commitments and ESG reporting. For customers who care about supply chain emissions—and more of them do each year—it is also a differentiator worth surfacing.

IX. What This Actually Tells Us

A few things stand out from this case that are not obvious from reading the headline numbers.

First, the economics work before full autonomy arrives. The 3PL is not betting on driver-out as the financial justification—it has already happened in the pilot data, at $32.30 versus $43.82 per stop, with a human safety driver still in the vehicle. The commercial case was built around that reality, not around a future driverless state. The autonomous upside is real but not load-bearing.

Second, the contract is doing a lot of work. The combination of the month-12 pricing reset option, termination for volume, and no-fee exit from month 26 means the 3PL has committed to far less than a 60-month liability. The minimum real commitment—the period from which there is genuinely no exit—is roughly 12 months. Every provision after that gives the 3PL increasing flexibility. Most coverage of AV logistics deals focuses on the contract value; this one is more interesting for its optionality structure.

Third, the supply-side urgency was real. A 12-month manufacturing lead time and a filling order book meant this was not a decision the 3PL could table for six months while running another RFP. Several major retail and CPG companies were already in the queue. Pricing was expected to rise. The pilot had been running long enough to be confident in the data. The conditions for decision-making were unusually clear.

Fourth, and perhaps most practically useful for other operators: the SLA improvement in EE% turned out to matter more than the financial savings in client conversations. A 26-point gain in Dallas-Fort Worth is the kind of number that shows up in client delivery reports, not just internal P&Ls. Operators evaluating AV logistics partnerships should probably lead with service-level evidence rather than cost-per-stop math when talking to commercial clients.

There are real limits to this analysis. The financial projections rest on a 2% annual courier inflation assumption that history suggests is optimistic for the carrier. If inflation runs higher—and the 2021–2022 experience suggests it can—the savings widen considerably and the model understates the investment's value. And the California regulatory situation remains an open variable that could force route redesign in that market.

X. Conclusion

The 3PL's expansion of the AV startup's fleet is not a proof-of-concept or a research pilot. It is a 54-truck, three-market operating decision backed by real performance data and a contract negotiated to protect the buyer at every stage of the technology's maturation. The numbers work: $6.5M in savings over five years, positive economics from day one, and service-level improvements that clients can actually see in their delivery metrics.

What makes it worth writing about is not the headline outcome but the decision architecture. The investment works because the contract is structured to absorb technology risk, because the pilot evidence was sufficiently robust to underwrite the expansion, and because the supply-side constraints created a real cost to waiting. That combination—solid pilot data, favorable contract terms, and a closing window—is a template that other logistics operators can evaluate against their own situations.

Autonomous middle-mile logistics is no longer a future state. For operators running fixed-route, repetitive D2DDU or hub-to-spoke networks, it is a present-tense financial decision. The remaining question is not whether the technology works. It is whether your contract protects you well enough to take the bet.

Acknowledgment

The author thanks the 3PL's Transportation and Operations teams for operational data and support, and the AV startup's commercial team for transparency throughout the pilot evaluation process.

References

[1] AV Startup, "Autonomous Vehicle Fleet Operations – Commercial Overview," Technical Briefing, 2023.
[2] 3PL Logistics Company, "Capital Investment Proposal – AV Startup Expansion Program," Internal CIP, Aug. 2023.
[3] 3PL Logistics Company, "AV Startup Expansion Executive Summary," Internal Memo, 2023.
[4] SAE International, "Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles," SAE J3016, 2021.
[5] University of California San Diego, Jacobs School of Engineering, "Autonomous Trucking and Fuel Economy," 2022.
[6] National Renewable Energy Laboratory, "Fuel Savings from Adaptive Cruise Control in Commercial Trucking," NREL Technical Report, 2021.
[7] U.S. Bureau of Labor Statistics, "Producer Price Index – Truck Transportation of Freight," Annual Series 2016–2022.
[8] FleetOwner, "AV Technology Fuel Efficiency in Heavy-Duty Fleets," 2022.
[9] U.S. DOE and EPA, "Fuel Economy Guide for Commercial Vehicles," 2023.
[10] BC Ministry of Transportation and Infrastructure, "Commercial Vehicle Eco-Driving Guidelines," 2020.

Author

Ashvin Kulkarni – Senior Technical Product Manager, Collaborative Innovations. Led the pilot evaluation, capital investment analysis, expansion planning, and on-site route implementation across all three markets (DFW, ONT, SCK).

Correspondence: Ashvin Kulkarni, Collaborative Innovations.