PITON Global

EXECUTIVE WHITE PAPER

AUTONOMOUS VEHICLE REMOTE OPERATIONS · PHILIPPINES

The availability standard: the executive economics of autonomous vehicle remote-operations outsourcing to the Philippines.



An analysis of why requests answered is a volume vanity metric, how guidance accuracy and minimal-risk-condition discipline — never response volume — decide the true cost of a remote-assistance desk once wrong guidance, vehicles left in live lanes, rescue dispatches and lost in-service hours are counted, and the vendor-selection discipline that guides the vehicle right and brings it to a minimal-risk condition before anything else.

AUTHORS

John Maczynski — Chief Executive Officer
Ralf Ellspermann — Chief Strategy Officer

MANILA · ESTABLISHED 2001

AUTONOMOUS VEHICLE REMOTE OPERATIONS

Contents

– About the authors 03

01 Executive summary 04

02 The volume mirage: requests answered versus vehicles returned to service 05

03 The guidance contract: triage to minimal-risk condition, guide the vehicle accurately, return to autonomy clean 06

04 Cost and performance: cost-per-correctly-guided-request economics and benchmarks 07

05 Sourcing discipline: the 7-Step Vendor Vetting Framework, applied to AV remote operations 08

06 Case study: a 60-seat robotaxi remote-assistance desk re-based on guidance accuracy 09

07 The executive playbook: governance for the first 180 days 10

– Methodology & notes 11

ABOUT THIS SERIES

Every AV remote-operations proposal is priced per request answered, because volume is easy to count. This volume is about the two outcomes a driverless fleet's economics actually depend on and a cheap contract quietly trades away: the request guided so the vehicle returns to service with its rider or load, and the safety-critical event brought to a minimal-risk condition — stopped safely, out of the live lane, escalated — before anything else. The gap between answering a request and guiding it safely is where lost in-service hours, wrong guidance, and lane-blocking exposure live. The full series lives at piton-global.com.

THE PITON-GLOBAL LEADERSHIP TEAM

About the authors



John Maczynski

Chief Executive Officer

John has bought remote-operations capacity for three decades of the outsourcing era, and learned that the fast desk that answers every request and guides the vehicle wrong is the most expensive one a fleet operator can hire — because wrong guidance can put a robotaxi into oncoming traffic, an unanswered request is a rider stopped in a live lane, and a missed minimal-risk stop is an incident no rate card can price. His standing question ignores the requests-per-hour report: of the requests you answered, how many put the vehicle back in service — with every safety event brought to a minimal-risk condition first? Applied at PITON-Global, that question is what clients get on every AV remote-operations sourcing decision — personally, with no account layer in between.

j.maczynski@piton-global.com
+1 402-598-8740



Ralf Ellspermann

Chief Strategy Officer

Ralf has spent twenty-five years running Philippine operations-support floors, and he judges a remote-assistance desk by its guidance accuracy and its minimal-risk-condition discipline, not its requests-per-shift: whether the vehicle was guided back into service and every safety event stopped safely, or cleared off the queue and left on the shoulder. The guidance contract in Section 3 codifies what he learned rebuilding desks that closed requests fast and left vehicles stranded. He now reads AV remote-operations providers the way he once ran them — from the availability log and the minimal-risk audit backward.

r.ellspermann@piton-global.com
Manila, Philippines

PITON Global

ABOUT PITON-GLOBAL

PITON-Global is a buyer-side advisory practice, Manila-based since 2001, that has never operated a delivery floor or taken a placement commission. Its stock-in-trade is knowing the Philippine provider market at the operating level — 1,000+ mapped operations, re-verified continuously, including each remote-operations team's real guidance accuracy and minimal-risk-condition record rather than its requests-per-shift headline — and converting that knowledge into a shortlist of six to ten genuinely qualified partners per engagement. Buyers from Tripadvisor, Garmin, TSYS, HealthPartners, Yetipay, Workrise, and Norman Disney & Young have run their searches through the practice.

SECTION 01

Executive summary

−56%
Fully-loaded cost per correctly-guided request vs. US onshore

96%+
Guidance accuracy on vetted teams, up from 78–86%

100%
Safety-critical events brought to minimal-risk condition to protocol, vetted teams

5.9×
First-year ROI, reconstructed case (Section 6)

An AV remote-operations engagement is bought on the price per request and paid for on the vehicles it fails to return to service — and, once in a while, on the minimal-risk stop it fails to honor. The vendor reports millions of requests answered at a fraction of the onshore cost; the fleet operator finds the requests cleared off the queue but the vehicle still on the shoulder, so rescue drivers keep rolling; the wrong path confirmations that send a robotaxi into a closed lane or a bus stop; the events handled as ordinary requests when protocol required an immediate minimal-risk stop; and an in-service-availability number sliding while the request count looked healthy. The desk cleared its queue and the operator inherited an availability problem, a guidance-quality problem, and a public-road safety exposure it never priced. The gap between answering a request and guiding it — safely — is where the real cost of cheap remote assistance hides, because an unguided or unsafe request is not throughput delivered; it is a vehicle that is not carrying anyone, guidance that made things worse, and, at the extreme, a hazard in a live lane.

The Philippines — English-fluent, technically trainable, with a deep 24/7 monitoring and dispatch pool and follow-the-sun coverage of every US metro — is where AV remote operations are delivered with the most guidance-and-safety discipline, because that pool, properly trained and console-certified, can run the

accurate path confirmation, remote-driving rigor, and minimal-risk protocol that a public-road fleet demands. Five findings:

1. **Requests answered is the industry's most flattering number.** Any desk can acknowledge a pull-over request and close it. Vetted teams are measured on outcome: guidance accuracy, in-service availability restored, and flawless minimal-risk-condition discipline. Weak providers quote a low per-request price the operator repays in stranded vehicles, bad guidance, and public-road risk. Section 2 shows the gap.

2. **The vehicle returned to service, not the request answered, is the product.** Guidance that protects fleet economics comes from reading the scene correctly — the double-parked truck, the hand-directed intersection, the blocked pick-up zone — and confirming a path the vehicle can drive, not from clearing a queue. Desks that guide accurately close the request once; desks that answer loosely generate repeat requests, rescue rolls, and lost trips that cost far more than the request.

3. **Minimal-risk-condition discipline is the one call where a single miss is catastrophic.** A vehicle with a rider inside, an emergency vehicle approaching, a sensor fault of uncertain severity must be brought to a minimal-risk condition — stopped safely, out of the live lane, escalated — before any availability consideration. Vetted teams hold minimal-risk discipline at 100% to protocol precisely because the cost of the one missed stop dwarfs every efficiency the contract could buy.

4. **Lost in-service hours and rescue dispatches, not the per-request rate, are what cost.** A request guided wrong costs the repeat request, the stranded rider, the rescue driver, and the lost fares or loads — many times the per-request saving. Desks that guide accurately cut repeats and protect availability, and that protected availability — not the offshore rate alone — is where the case study's return is built.

5. **Contract on the vehicle returned to service, not the request answered.** The contracting failure of AV remote operations is paying per request and hoping vehicles get back on the road and minimal-risk stops are honored. Vetted deals price against guidance-accuracy floors, availability-restoration targets, and a non-negotiable 100% minimal-risk standard, making guidance and safety the desk's problem, which is where they belong.

The intended reader is a VP of autonomy operations, head of fleet operations, or COO at a robotaxi, autonomous-trucking, or ADAS-fleet operator buying remote assistance per request and paying for it in stranded vehicles, bad guidance, and public-road exposure. Sections 2–4 separate the desks that guide vehicles safely from the desks that clear requests; Sections 5–7, the method and the proof. This paper

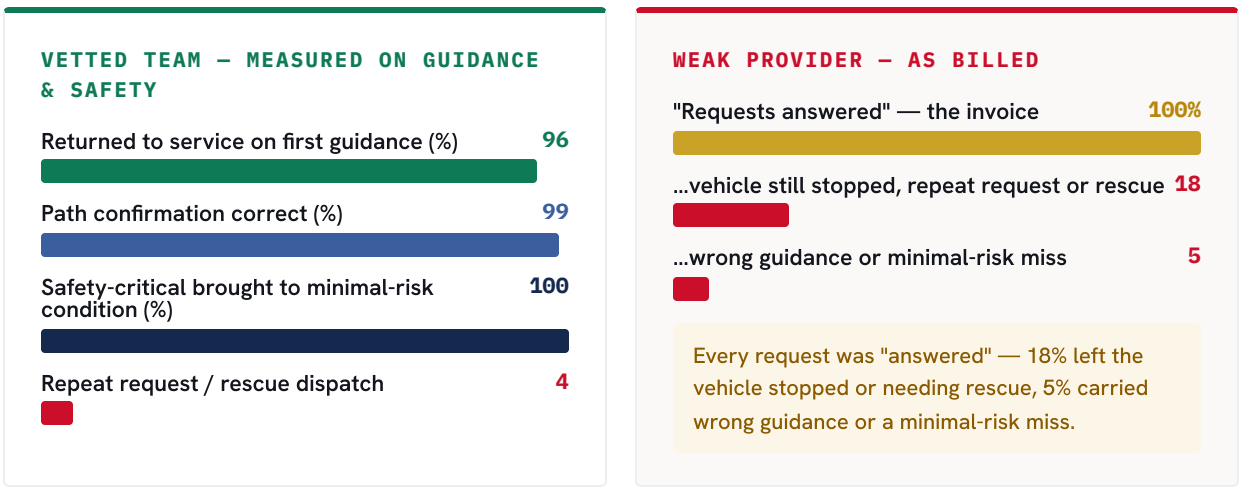
addresses human remote operations around autonomous vehicles; it is not engineering, functional-safety-certification, or regulatory advice, and minimal-risk protocols and remote-driving authority are always operated to the fleet operator's own engineering-approved specifications and jurisdictional permits.

SECTION 02

The volume mirage: requests answered versus vehicles returned to service

Figure 1 shows one week of a robotaxi fleet's remote-assistance requests two ways: as the requests-answered report a weak vendor bills, and as the guidance-and-safety reality the availability log, the repeat-request tracker, and the minimal-risk audit experience. The requests-answered number is engineered to look like output while lost in-service hours, repeat requests, and mis-handled safety events accumulate underneath.

FIG. 1 — ONE WEEK OF REMOTE-ASSIST REQUESTS: GUIDANCE-AND-SAFETY QUALITY VS. REPORTED VOLUME



Left: PITON-Global engagement instrumentation, 2025–26, guidance and minimal-risk audited. Right: composite of pre-engagement vendor reporting reviewed in diligence.

The economics follow the guidance and the safety, not the request count. On the left, ninety-six percent of requests return the vehicle to service on the first guidance and every safety-critical event is brought to a minimal-risk condition, so availability holds and no lane is blocked; on the right, a cheap per-request price leaves eighteen percent stopped or rescued and errs on guidance or minimal-risk for five, paying in lost trips, rescue rolls, and public-road exposure at once. A low per-request price that leaves vehicles on the

shoulder is not a saving; it is an availability treadmill and a risk the operator carries. The guidance-and-safety audit — guidance accuracy, availability restored, minimal-risk compliance — is the cheapest due-diligence instrument in this paper.

SECTION 03

The guidance contract: triage to minimal-risk condition, guide the vehicle accurately, return to autonomy clean

Figure 2 lays out the three clauses of a working AV remote-operations contract — the machinery a floor read verifies before any per-request price is believed.

FIG. 2 — THE GUIDANCE CONTRACT, AS OPERATED ON VETTED TEAMS

TRIAGE TO MINIMAL-RISK CONDITION

Every request is screened for a safety-critical signal — a rider on board with a stalled vehicle in a live lane, an approaching emergency vehicle, a sensor or connectivity fault of uncertain severity — and any such event is brought to the fleet operator's engineering-approved minimal-risk condition and escalated in seconds, never queued. Safety is the first gate, before availability ever enters the picture.

GUIDE THE VEHICLE ACCURATELY

The scene — a double-parked truck, a hand-directed intersection, a blocked pick-up zone, a construction detour — is read correctly and the right path confirmation or low-speed remote-driving action taken within approved authority, with anything beyond that authority escalated to the safety team rather than improvised. Correct guidance is guidance the vehicle can drive and the fleet does not request again at the same spot.

RETURN TO AUTONOMY CLEAN

The vehicle is handed back to autonomy with the request fully documented — scene, guidance, telemetry, rider communication — so the availability and repeat-request data are trustworthy and the autonomy team gets the labeled scenario it needs, with a guidance report the operator can read. A clean return is the outcome the availability line and the driver roadmap both feel.

↓ yields ↓

THE GUIDED, SAFE REQUEST

96%+ GUIDANCE ACCURACY **100% MINIMAL-RISK CONDITION** **AVAILABILITY RESTORED**

Requests are guided back to service, every safety event is brought to a minimal-risk condition, and every hand-back is documented clean — so lost in-service hours fall, repeat requests drop, and no lane is blocked. The desk gets cheaper per guided request and safer at the same time, which is the entire promise of outsourcing AV remote operations to people who measure guidance and safety, not requests answered.

Operating parameters from PITON-Global AV remote-operations engagements, 2025–26. Minimal-risk protocol and remote-driving authority are operated to each fleet operator's engineering-approved specifications and permits.

In diligence, the contract is tested from the fleet's and the road's side: pull answered requests and check the repeat-request and rescue-dispatch trackers for vehicles that stayed stopped; re-audit a guidance sample against the approved authority and the scene; and stress-test minimal-risk discipline with scripted event scenarios in simulation — because a desk that misses one minimal-risk stop fails the whole engagement, whatever its efficiency stats say. Perhaps one team in ten survives it cleanly. Steps 3 and 5 of the framework exist to find that one before the contract does — because everything in Table 3 of Section 6 depends on guided, safe requests, not requests answered.

SECTION 04

Cost and performance: cost-per-correctly-guided-request economics and benchmarks

TABLE 1 — FULLY-LOADED COST PER CORRECTLY-GUIDED REQUEST: THREE OPERATING MODELS (2026)

OPERATING MODEL	COST / REQUEST	VS. US BASELINE
US onshore in-house remote-assistance desk	$16.40	baseline
Low-cost offshore request mill (per-request)	$11.10	−32%
Vetted Manila team — guided, safe, clean	$7.22	−56%

Basis: PITON-Global engagement pricing 2025–26; cost per correctly-guided request = fully-loaded operating cost divided by requests that returned the vehicle to service, clean, and safely (repeat, rescued, and mis-guided excluded); the request mill looks cheap per request but lost trips, rescue dispatches, and public-road exposure lift its true cost per guided request; US comparator fully-loaded in-house cost.

TABLE 2 — AV REMOTE-OPERATIONS BENCHMARKS, VETTED VS. BASELINE (2025–26)

METRIC	VETTED TOP TIER	BASELINE	EXECUTIVE READING
Guidance accuracy	96%+	78–86%	The only support number availability follows
Minimal-risk-condition compliance	100%	91–98%	Every event stopped safely, or one missed
Rescue-dispatch rate vs. baseline	−64%	baseline	"Guide the vehicle" — held or not
In-service availability contribution	+2.8 pts	baseline	Vehicles carrying riders, or waiting
Remote-assist response time	−91%	baseline	Six seconds, or three minutes

Source: PITON-Global AV remote-operations engagement data 2025–26; baseline = typical pre-engagement operations billed on requests answered.

The strategic read: the request mill's –32% per request is a false economy — lost trips, rescue dispatches, and public-road exposure lift its true cost per guided request far above the per-request gap, and a single missed minimal-risk stop is a cost no rate recovers. The vetted Manila team's –56% is real because it is measured where the value is: the vehicle returned to service, the guidance correct, the event brought to a minimal-risk condition. Buy on requests answered and you buy the right column of Figure 1; buy on cost-per-correctly-guided-request and you buy guidance and safety — the same discipline every volume in this series keeps arriving at from a different door.

SECTION 05

Sourcing discipline: the 7-Step Vendor Vetting Framework, applied to AV remote operations

AV remote-operations diligence adds two hard requirements to every step: a per-request price is real only when the guidance accuracy behind it survives a repeat-request and rescue-dispatch audit, and only when the minimal-risk protocol behind it survives scripted simulation stress tests. The funnel below is representative of a 40–120-seat AV remote-operations search.

FIG. 3 — THE 7-STEP FRAMEWORK AS AN AV REMOTE-OPS FUNNEL

	STEP	WHAT IT ELIMINATES IN AN AV REMOTE-OPS SEARCH	FIELD REMAINING
1	**Requirements definition**	Volume-first scoping: fixes the guidance-accuracy floor, the availability-restoration target, the rescue-dispatch ceiling, the response-time SLA, and the non-negotiable 100% minimal-risk requirement by event type	1,000+
2	**Market mapping**	Request mills: operations that bill request volume and treat stranded vehicles and mis-guidance as the client's problem, and generalist floors with no AV, remote-assist console, or minimal-risk-protocol depth	~31
3	**Capability screening**	Teams without the machinery: audited guidance-accuracy history, remote-assist console and fleet-platform certification, a minimal-risk protocol with drilled escalation, latency-budgeted connectivity, and coverage across your operating domains and metros	~10
4	**Compliance & security audit**	Weak governance: camera, lidar and rider-PII protection, access controls on remote-assist and fleet systems, incident-reporting standards a regulator will accept, dual-site continuity, SOC 2 scope covering the ops stack	~6
5	**Operational diligence**	The guidance-and-safety read: audit repeat-request and rescue trackers for vehicles left stopped; re-check a guidance sample against approved authority and the scene; run scripted minimal-risk drills in simulation; interview on a complex emergency-vehicle escalation	4–6
6	**Competitive RFP & negotiation**	Per-request pricing: finalists bid against guidance-accuracy floors, availability-restoration targets, the response-time SLA, and the 100% minimal-risk standard — all carrying credits for stranded vehicles and any minimal-risk miss	2–3
7	**Transition & launch governance**	Cold cutover: the team certifies on your remote-assist console, operating domains, and minimal-risk protocol in shadow mode before it owns the live fleet, with the guidance floor, response SLA, and minimal-risk discipline certified before steady state, PITON-Global embedded buyer-side	1

Field-remaining counts represent a 40–120-seat AV remote-operations search, PITON-Global engagements 2024–26.

The standing guarantees hold: **vendor neutrality** — PITON-Global operates no remote-operations floors, answers no requests, and takes no placement economics, so the shortlist reflects audited guidance and minimal-risk performance, not inventory; **right-sizing** — teams for whom your operation is a flagship account. Six to ten genuinely qualified providers, delivered within 48–72 hours, free to the buyer, competing through a fully transparent, fair, comprehensive, and highly competitive RFP.

SECTION 06 · ANONYMIZED CASE STUDY

Re-basing the operation on guidance accuracy: a 60-seat robotaxi remote-assistance desk rebuild, reconstructed

The case engagement is a US robotaxi operator — running roughly 540 L4 vehicles across two metros and generating about 2,100,000 remote-assistance requests a year (pull-over requests, path confirmations, blocked pick-up zones, emergency-vehicle interactions, connectivity faults) — that had offshored remote assistance to a per-request mill. The request price was excellent; the guidance and safety were not: guidance accuracy sat near 81%, vehicles stayed stopped and rescue drivers rolled, response time had crept past two minutes, and a simulation drill had found an emergency-vehicle approach handled as a routine request when protocol required an immediate minimal-risk stop. The VP of autonomy operations wanted the arbitrage without the stranded riders and the public-road exposure. Identifying details are anonymized under NDA.

SELECTION (WEEKS 0–8, GUIDANCE + MINIMAL-RISK READS)

Requirements fixed a 96% guidance-accuracy floor, an availability-restoration target, a rescue-dispatch ceiling, a sub-6-second response SLA, and a non-negotiable 100% minimal-risk requirement. Mapping produced 31 claimants; the machinery screen — audited guidance accuracy, remote-assist console certification, drilled minimal-risk protocol, latency-budgeted connectivity, metro coverage — left 10; the compliance-and-continuity audit left 6. Repeat-request audits, guidance re-checks, and scripted minimal-risk drills in simulation shortlisted 4. A Manila–Clark dual-site team won priced against guidance and minimal-risk, with both written into the reporting spec.

TRANSITION (MONTHS 2–7, SHADOW MODE FIRST)

The team certified on the operator's remote-assist console, operating domains, and minimal-risk protocol in shadow mode before owning the live fleet; the guidance floor, response SLA, and minimal-risk discipline were certified in month 3, and the minimal-risk drill ran weekly from day one. The guidance-and-QA loop fed operator coaching continuously. By month 7 guidance accuracy reached 96.4%, rescue dispatches fell 64%, response time dropped from over two minutes to under six seconds, in-service availability rose 2.8 points to 99.1%, and minimal-risk discipline held at 100% across every drill. PITON-Global chaired the guidance-and-safety review board to steady state.

TABLE 3 — CASE FIRST-YEAR VALUE BUILD-UP (US$)

VALUE COMPONENT	YEAR 1	BASIS
Labor arbitrage on the 60-seat desk	**$1.72M**	112.6k productive hrs × ($25.30 − $10.00)
Ride revenue recovered (higher availability)	**$1.34M**	In-service vehicle-hours restored at guidance floor
Rescue-dispatch & repeat-request reduction	**$0.46M**	Rescue drivers, tow rolls and rider re-routing removed
Safety & incident exposure avoided	**$0.31M**	Minimal-risk miss and lane-blocking exposure removed
Gross value created	**$3.83M**	
Less: engagement & transition costs	**($0.65M)**	Console certification, minimal-risk-drill build, latency and failover integration; advisory fee $0 (provider-paid model)
Net value ÷ cost = first-year ROI	**5.9×**	$3.83M ÷ $0.65M

Figures rounded; guidance, availability, and minimal-risk effects audit-verified. Method in Methodology & notes.

Forty-five percent arbitrage, fifty-five percent from recovered ride revenue, reduced rescue dispatches, and avoided safety exposure — a heavier value-side split than most of the series, earned because a stopped robotaxi with a rider inside is a more expensive asset than most desks support. Every line traces to a diligence step: the guidance floor and minimal-risk standard to Step 1, the guidance-and-safety pricing to Step 6, the minimal-risk-drill protocol to Step 3's contract. The VP's line at the year-one review: "The request price barely moved. What moved is that the cars stopped sitting on the shoulder — and I stopped worrying about the emergency vehicle we'd wave through."

SECTION 07

The executive playbook: governance for the first 180 days

An AV remote-operations transition succeeds when the team is measured and paid on guidance accuracy, response time, and minimal-risk discipline — client in command throughout.

DAYS 0–30

Define guidance & minimal risk, contract on them

Set the guidance-accuracy floor, the availability-restoration target, the rescue-dispatch ceiling, the response-time SLA, and the non-negotiable 100% minimal-risk requirement with autonomy-operations and safety leadership before launch, and document the minimal-risk and escalation protocol by event type and operating domain. Contract against guidance and safety — never per request alone — with floors carrying credits and any minimal-risk miss a defined breach. Stand up the repeat-request audit and the simulation minimal-risk drill as the program's arbiters of truth.

DAYS 30–90

Shadow, certify, drill the minimal-risk protocol

Certify the team on your remote-assist console, operating domains, and minimal-risk protocol in shadow mode before it owns the live fleet; certify the guidance floor and response SLA and run the minimal-risk drill to 100%. Switch on the guidance-and-QA loop from day one. Publish the operations dashboard: guidance accuracy, repeat-request rate, rescue dispatches per thousand miles, remote-assist response time, in-service availability, minimal-risk compliance.

DAYS 90–180

Own the fleet, certify steady state

Hand the live fleet to the team only after it holds the guidance floor, the response SLA, and 100% minimal-risk discipline. Review the board monthly with the repeat-request and minimal-risk-drill results on the table. Certify steady state on two months at floor including a new-metro launch or holiday-demand surge; hand over on documented criteria; keep the repeat-request audit and weekly minimal-risk drill standing.

THE FIVE FAILURE MODES WE ARE HIRED TO PREVENT

Paying per request while vehicles sit on the shoulder · clearing requests instead of guiding vehicles · treating an emergency-vehicle event like an ordinary request · improvising remote driving beyond approved authority · owning the fleet before guidance and minimal risk are proven. All five are settled at selection and contract.

The AV remote-operations argument, in one sentence: an answered request is activity and a guided, safe request is a vehicle back in service and a hazard that never reached a live lane, and a vetted Philippine team delivers –56% cost per guided request with 96%+ guidance accuracy, rescue dispatches cut by nearly two-thirds, sub-6-second response, and 100% minimal-risk discipline — for buyers who audit guidance and drill minimal risk instead of counting requests answered, and compete the finalists through a fully transparent, fair, comprehensive, and highly competitive RFP. Anyone can answer a request. We make sure you buy the desk that guides the vehicle and brings the event to a minimal-risk condition.

METHODOLOGY & NOTES

How the numbers were built

Guidance and safety figures (Fig. 1, Table 2) derive from engagement instrumentation in PITON-Global AV remote-operations engagements, 2025–26. Guidance accuracy is measured against the repeat-request and rescue-dispatch trackers; minimal-risk compliance by scripted simulation drills; rescue-dispatch rate, remote-assist response time, and in-service-availability contribution against the client's pre-engagement baseline. The weak-provider column is a composite of pre-engagement vendor reporting reviewed in diligence, shown for contrast.

Cost-per-correctly-guided-request figures (Table 1) are fully loaded — wages, benefits and statutory charges, facilities, technology and remote-assist console/connectivity tooling, management and QA overhead including minimal-risk-drill leads, and attrition-driven recruiting — divided by requests that returned the vehicle to service, clean, and safely, with repeat, rescued, or mis-guided requests excluded from the denominator; US comparators reflect fully-loaded in-house cost on the same basis. A correctly-guided request returns the vehicle to autonomous service, is documented clean, and — where safety-critical — is brought to the operator's engineering-approved minimal-risk condition.

The case study is anonymized under NDA; volumes and dates are generalized, figures rounded. Guidance, availability, and minimal-risk effects are audit-verified; the ride-revenue-recovered line reflects restored in-service vehicle-hours at the operator's average fare yield, and the exposure-avoided line is a conservative estimate. ROI = net first-year value ÷ all engagement-related client costs. PITON-Global's advisory fee is provider-paid and appears at $0 in the client cost base; neutrality safeguards are documented at piton-global.com.

Scope and market context. This paper concerns human remote operations around autonomous vehicles — fleet monitoring, remote assistance, low-speed remote driving, and minimal-risk triage — and is not engineering, functional-safety-certification (ISO 26262 / UL 4600), or regulatory advice; minimal-risk protocols and remote-driving authority are always operated to the fleet operator's own engineering-approved specifications and jurisdictional permits. Context draws on IBPAP reporting and PITON-Global's proprietary provider mapping (1,000+ operations; guidance and minimal-risk performance re-verified August 2026). Estimates and projections are PITON-Global analysis and should be cited as such. The full series is available at piton-global.com.



Would your provider's request price survive a guidance-and-minimal-risk audit?

Thirty minutes with John will tell you. Free to buyers, strictly vendor-neutral — a guidance-and-minimal-risk-verified shortlist within two to three business days.

j.maczynski@piton-global.com
+1 402-598-8740

© 2026 PITON-GLOBAL, INC. · EXECUTIVE WHITE PAPER SERIES · AUTONOMOUS VEHICLE REMOTE OPERATIONS · PHILIPPINES. This publication is provided for general information and is not engineering, functional-safety-certification, or regulatory advice. Figures identified as estimates or models are PITON-Global analysis. Reproduction with attribution is permitted.