



**MIND THE
AI GAP**

Transport operations 2026:

How leading operators
address challenges shaping
the next two years





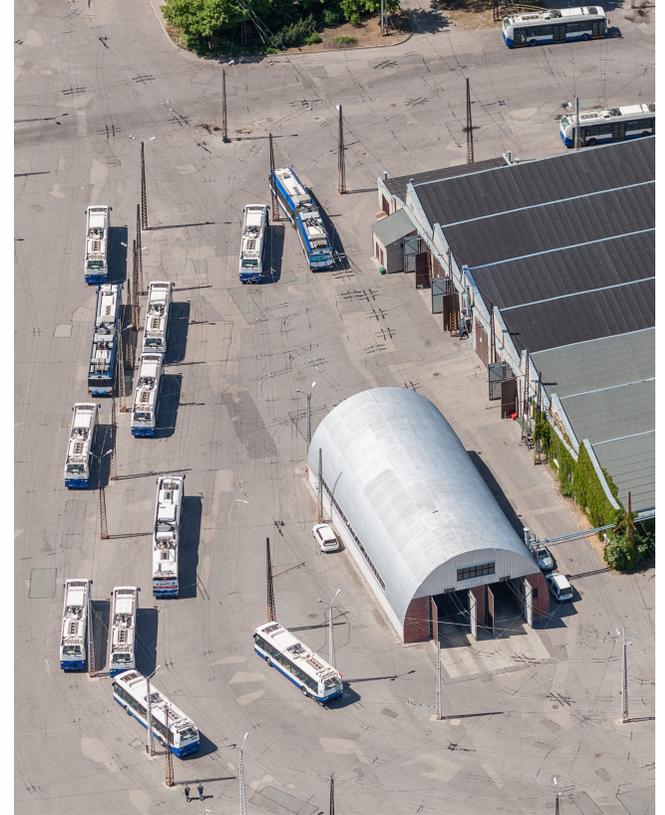
Keeping services running to schedule has always been key. But nowadays this is much harder. **The pressure is coming from multiple directions.**

For starters, disruption has increased. For instance, UK drivers lost **62** hours to congestion in 2024, costing the economy £7.7 billion. Plus, average speeds on major roads dropped **3.8%** in just two years, with delays per mile up 26%.

At the same time, passenger expectations around journey planning and service updates continue to rise. Plus, travel patterns have changed with hybrid working, stretching planning models built around traditional peak and off-peak demand.

And then there's the economics.

Fuel, energy, maintenance, and recruitment costs keep rising while margins in franchised areas narrow. Add funding uncertainty into the mix, and planning and investment decisions become increasingly difficult to make with confidence.



So where does that leave transport operators in 2026?

This guide reviews where pressure is building most intensely across transport operations, and more importantly, how the operators managing these challenges most effectively are actually doing it.

7 operational pressures that keep resurfacing

Speak to almost any transport operator, and you'll hear the same issues coming up again and again. And they cut across all organisations, including bus and rail, geography, and organisational structure. **The context changes. The pressure doesn't.**

Here's what makes 2026 different: individually, none of these pressures are new. What's changed is how often they now overlap, and how little room operators have left to absorb them. When one pressure hits, you can manage it. When three or four hit simultaneously, the system starts to strain.

01. Running services to schedule and keeping customers informed

Keeping services running to schedule (*and keeping customers informed when they don't*) remains one of the most persistent challenges across transport. Rising congestion continues to affect punctuality and journey times. Disruption is more frequent. At the same time, customer expectations around journey planning, ticketing, and service updates have risen. Passengers expect timely, accurate information, not explanations afterwards.

Travel trends have also changed. Hybrid working has reshaped demand, but many planning models still rely on historic peak and off-peak assumptions that no longer reflect how services are actually used.

02. Driver shortages and people pressure across operations

Driver shortages persist despite improvement in some areas. Bus vacancies dropped to **3.4%**, but coach operators face a worsening crisis at 12.4%. More concerning: per-mile operating costs rose 17% year-on-year as operators competed for scarce talent.

In many organisations, this is treated as a structural reality to live with. The focus moves to marginal recruitment improvements rather than reducing pressure on the system itself. Alongside this, the skills required to run modern operations are changing. Technology and data capability gaps are often addressed tactically through training courses, rather than as a deeper organisational capability challenge.



03.

Safety performance that remains under scrutiny

Safety performance deteriorated in 2024-25. Preventable railway fatalities rose [26%](#), while signals passed at danger increased to 305, which is the highest in years.

Too often, action follows incidents rather than preventing them. Early indicators exist, but they are not always visible or acted on in time. Safety is therefore reviewed retrospectively, rather than managed proactively.

04.

Rising costs and tightening margins

The cost of running a service continues to rise, particularly in franchised areas where margins are narrowing at the same time. For instance, bus fares have risen [93.5%](#) since January 2015, including a 22% increase in the past year alone. For haulage operators, the situation is even more acute: pre-tax profit for the top 100 haulage companies dropped to just [1.58%](#) in 2024, down from 2.6% the previous year. Recruitment costs add further pressure.

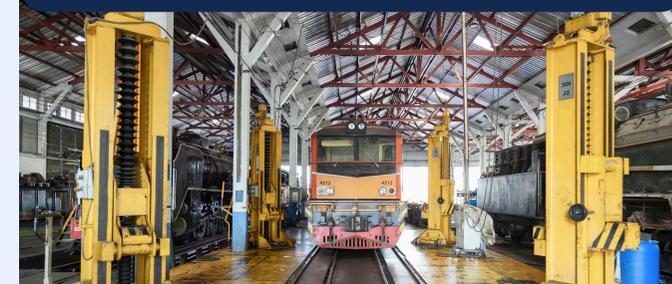
Longer-term funding uncertainty, including programmes such as BSIPs, makes it harder to plan and invest with confidence. For many leadership teams, this leaves a narrowing window to improve efficiency without impacting service quality.

05.

Decarbonisation and long-term asset pressure

Decarbonisation introduces a different kind of challenge. Moving to electric fleets requires *significant* capital investment, alongside the infrastructure needed to support them. In the meantime, many operators are managing mixed fleets, adding operational complexity rather than removing it. Ageing vehicles, depots, and facilities often require major investment simply to remain viable.

These pressures sit alongside day-to-day operational demands, competing for the same attention and budget.

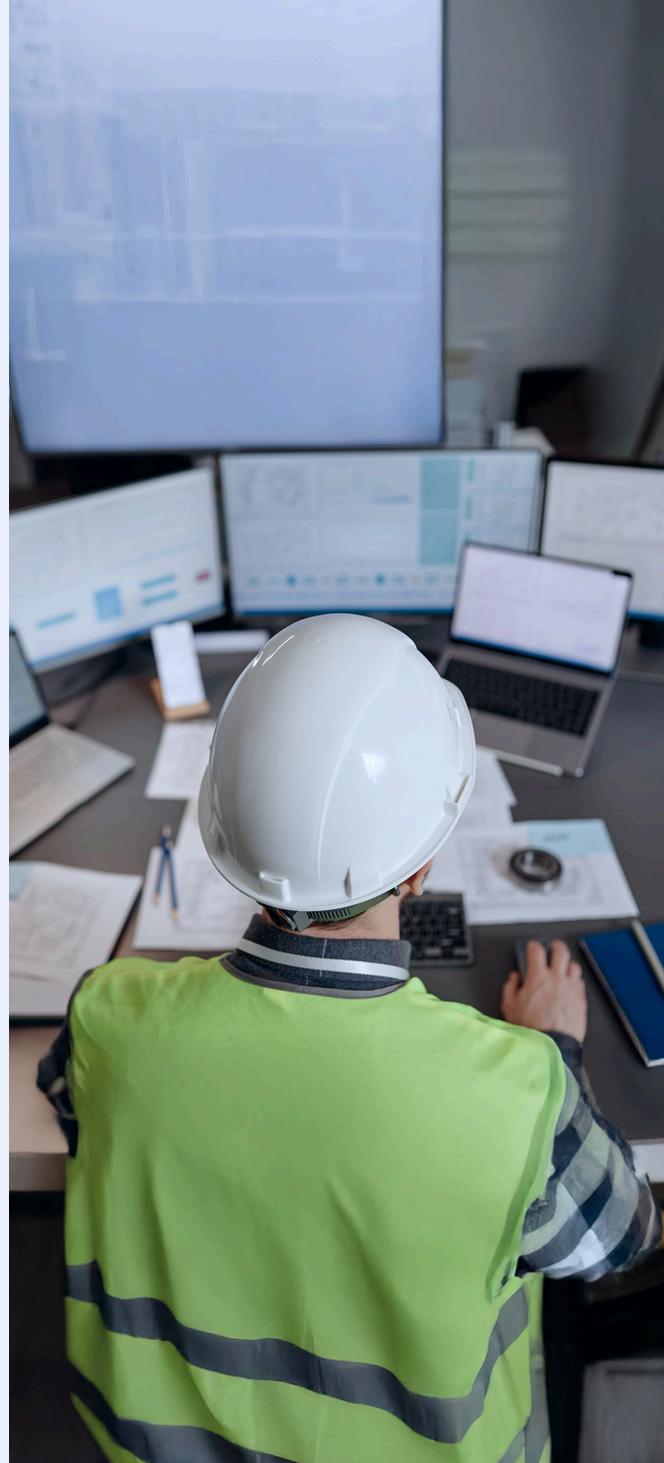


06.

Technology debt and limited access to real-time insight

Technology is intended to help simplify operations. In reality, it often adds to the burden. Many operators carry significant technology debt, with legacy systems built on older ones. Integration is difficult. Cybersecurity risk increases. Plus, access to real-time data and predictive analysis remains limited, even where data exists.

Point solutions across ticketing, fleet management, safety, scheduling, and customer information frequently operate in isolation, making coordinated decision-making harder than it needs to be.



07.

Industrial action and policy uncertainty

Industrial action continues to cause disruption, particularly on the rail network. Regulatory and policy uncertainty adds further complexity. In the UK, bus operators face ongoing uncertainty around franchising, while rail continues its transition toward Great British Railways.

Many of the pressures facing transport today are well understood. What is harder to acknowledge is how often organisations respond by reinforcing existing ways of working.

Where familiar approaches are consuming more time and money

Across operations, the same approaches that worked for decades continue to dominate. **The problem?** The environment around them has fundamentally changed.

Here's what happens when pressure builds: instead of questioning whether current models still make sense, organisations double down on what they know. More planned maintenance. More training courses. More resources are thrown at the same processes. The assumption is that if the old approach isn't working, the answer must be to do more of it.

The result? A widening gap between effort and outcome:

Reactive maintenance in a world of real-time data

Maintenance models in transport have barely changed for decades, despite widespread deployment of telematics and real-time vehicle data. Vehicles are still maintained based on mileage or fixed schedules. Some are taken off the road when they don't need to be, while early warning signs on others are missed entirely. Breakdowns continue to occur in service, causing cancellations and disruptions that could have been avoided.

The data exists. The issue is the continued reliance on scheduled intervention rather than condition-led decision-making. Efficiency gains remain unrealised, and breakdown rates stay stubbornly high.

Resource allocation is built on old-school peak and off-peak assumptions

Planning and resource allocation remain heavily shaped by pre-Covid peak and off-peak patterns. Operators recognise that travel behaviour has changed. Hybrid working has flattened traditional peaks and created new demand patterns throughout the day. Yet optimisation is still driven by planning cycles and agreements built around historic assumptions.

The effects are increasingly visible:

- Empty buses at traditional peak times
- Overcrowding at new demand peaks
- Capacity is deployed where it is least needed

As behaviour continues to change, the mismatch between supply and demand grows, even when overall resource levels remain unchanged.

Annual route reviews in a different demand environment

Route performance is still typically reviewed once a year, or only after a major complaint. In contrast, consumer platforms like Uber and delivery services *continuously* adjust routes and capacity. Transport operators often have the same underlying data, but remain locked into annual planning cycles because “that’s how we’ve always done it”.

The impact is subtle but significant. Patronage declines gradually, with limited visibility into why. By the time a route is formally reviewed, the opportunity to intervene early has passed.

Point solution procurement and the integration problem

Technology procurement in transport often reinforces the

problem it is meant to solve. Operators buy traditional point solutions from traditional suppliers. Ticketing, safety systems, Wi-Fi, fleet management, and scheduling all sit on the same vehicle but operate as separate systems. Integration becomes a nightmare.

Organisations spend significant sums on technology that does not talk to each other. Data is fragmented. Insight is delayed. Operational teams are left stitching together information manually when decisions need to be made quickly.

Periodic training in a continuously changing environment

Training models have also struggled to keep pace. Drivers and operational staff are still put through periodic training courses, even as technology, customer

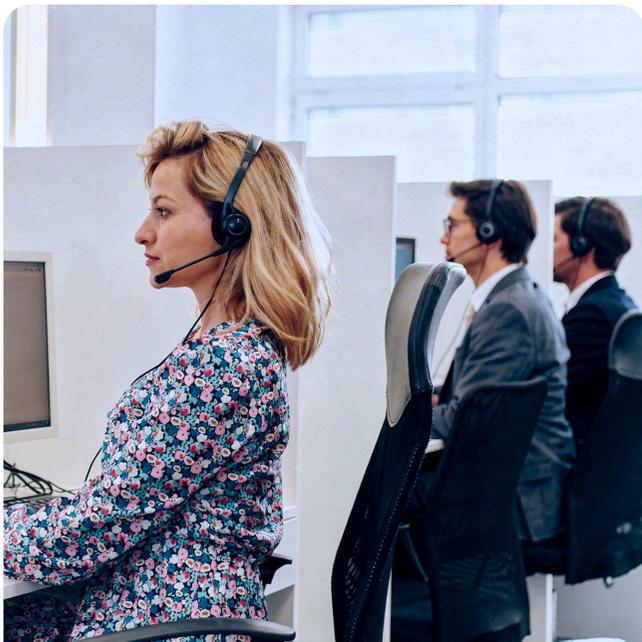
expectations, and safety requirements evolve rapidly. New systems are introduced, but adoption is inconsistent. Engagement drops. Knowledge gaps widen. What worked when change was slow no longer holds when systems, processes, and expectations are constantly shifting.

Hierarchical decision-making under operational pressure

In many organisations, front-line staff spot problems daily. Action, however, often depends on information moving up the chain before decisions are made. This slows response and adds pressure at exactly the point where speed matters. Decisions are delayed not by lack of awareness, but by how authority and information flow through the organisation.

Customer contact built around repetition

Customer contact remains heavily manual in many operations. Large teams handle repetitive queries like “Where’s my bus?” without a clear view of call patterns or root causes. This is costly, frustrating for staff, and inconsistent for customers. Voice agents and data analysis already help reduce repetitive demand while improving service quality. Yet many organisations continue to scale people rather than address the source of repetition.



Why familiar habits persist under operational pressure

What links these approaches is familiarity. They are embedded in long-standing processes, agreements, and ways of working that feel safe, even as their effectiveness declines.

The operators coping better than others are not immune to these pressures. The difference is where effort is concentrated. Rather than reinforcing manual processes, they focus on reducing friction in how work gets done. Existing data is used more consistently. Systems are joined up. And repetitive activity is reduced, so attention can move to issues that genuinely require judgment.

This is visible in day-to-day operations:

- Scheduling and maintenance decisions are taken earlier.
- Routine customer contact falls, allowing teams to deal with complex cases more effectively.
- Front-line staff are better supported with timely information, enabling decisions to be made closer to where problems emerge.

Here's how...

Where AI is already changing day-to-day operations

Issues that used to take days or weeks to surface are now visible immediately. Decisions that once required multiple meetings and sign-offs can happen in minutes. Maintenance that was scheduled months in advance can be triggered by the actual vehicle condition.

The examples below show what operators are doing today, using systems and data they already have in place:

Old way vs new way in Transport operations

Use case	Old way	New way
Predictive maintenance	Vehicle servicing based on mileage. Parts such as brake pads replaced during scheduled services, with vehicles off-road for days. Gearbox failures occur in service, causing breakdowns and cancellations.	Brake pad wear flagged four weeks earlier using telematics and braking patterns. Parts are pre-ordered and fitted during planned downtime. Gearbox failure detected around 10 days earlier through oil temperature analysis and repaired proactively, avoiding unplanned downtime and service impact
Optimised routing and logistics	Route 47 runs the same schedule it has for five years. Morning peak services run empty, while the 11am service is consistently overcrowded. Annual review scheduled in eight months.	Demand shift identified within two weeks. One morning peak bus is moved to the 10:45am slot. Operator approves and implements the change in the next schedule update. Capacity matches demand, passenger satisfaction improves, with no additional resource cost.
Demand forecasting and capacity planning	Planning teams review last year's data in February to build a September schedule. A weather event causes an unexpected demand spike. No spare capacity available. Passengers are left behind and complaints rise on social media.	Demand forecast continuously using weather, events, school calendars, and economic indicators. Demand spike predicted 48 hours ahead. Reserve vehicles are deployed. Capacity meets demand with minimal disruption.

[Continued on Page 9 >](#)

Use case	Old way	New way
Safety and surveillance	Driver records three harsh braking incidents over six months. Flagged at annual review and given generic refresher training. Driving behaviour does not change. A minor accident follows.	Harsh braking patterns identified within three weeks, clustering in wet conditions. Targeted coaching on wet weather driving triggered before the next shift. Driving behaviour improves and the accident is avoided.
Customer contact centre	Customer calls to ask "Where's my bus?" Agent checks the system and explains an eight-minute delay. Call takes four minutes. Multiple callers ask the same question about different routes. Call queues build to 15 minutes.	Voice agent answers instantly with real-time information: "Your 8:15 service is currently eight minutes delayed due to congestion on High Street, expected at your stop at 8:31." No queue time. Human agents focus on complex complaints and compensation. Call volume drops by around 50%.
Service disruption management	Accident closes a main road at 7:45am. Control room notified at 7:50am. Diversion options discussed. Drivers called individually. Website updated manually. Passengers find out when the service does not arrive. First diverted service appears at 8:20am.	Traffic anomaly detected at 7:46am using multiple data sources. Diversion routes generated automatically and pushed to driver apps. Passenger information screens and apps updated at the same time. Automated customer notifications sent. First diverted service runs by 7:52am, with passengers already informed.

What these examples have in common

None of these changes requires perfect systems or wholesale replacement of existing infrastructure.

They rely on:

- Data operators already generate
- Faster interpretation of that data
- Earlier intervention
- Better coordination across teams

Decisions move closer to where issues emerge. Disruption is addressed earlier. Resources are used more effectively. The change is less about new tools and more about when decisions are taken and who is able to act on them.

Here's an example...

Real AI use case:

AI voice agents in customer support

AI voice agents were implemented by a major US operator to support its customer service function, handling **800,000+ minutes of calls per month** across product lines.

What was implemented

The system now touches **95% of all incoming calls**, with **10% of complex support interactions fully automated end-to-end**. For the remaining calls, intelligent intake and basic troubleshooting save an **average of 2.5 minutes per call**.

Operational impact

Customers receive immediate, expert-level support for critical issues such as **product registration and order tracking**, often after

already spending **around 45 minutes attempting to resolve issues themselves**. This removes queue time and speeds up resolution at moments of highest frustration.

Why this worked

What sets this implementation apart is the **depth of customisation**, rather than generic chatbot deflection. The voice agents are trained on the operator's internal knowledge base, supported by continuous learning, and guide customers through step-by-step diagnostic processes.

Human operations teams **supervise and tune the system weekly**, keeping it aligned with real customer issues. The focus is on **resolution rather than call avoidance**.

Delivery timeframe

Deployment was completed in **six weeks**, showing that advanced customer support automation can be implemented quickly when built around real product knowledge and operational needs.



What separates operational leaders in 2026

By 2026, differences in operational performance will be clearer across the sector.

Those pulling ahead are combining long-standing operational knowledge with new capability to support decisions earlier in the operational cycle. Experience and judgement remain central, but they are applied more consistently and across a wider range of decisions.

Relying solely on experience is becoming harder to sustain as pressure increases. Operators that support experienced teams with better use of data respond faster, run leaner operations, and create space to focus on the issues that require attention.

Transport does not run on data alone. It runs on judgment. Depot managers understand shift patterns, local road conditions, seasonal variation, and behaviours that do not appear cleanly in systems. That knowledge remains essential. Earlier visibility and faster feedback allow it to be used more effectively across the organisation.

What this looks like in practice

In organisations pulling ahead, several patterns appear repeatedly:

- Decisions are taken earlier, using timely rather than perfect information
- Issues are addressed closer to where they emerge
- Experienced staff spend less time reacting and more time resolving
- Manual activity is reduced where it adds limited value

Judgment still drives everything. The difference is how early it can be applied and how many people can benefit from it.

Why partner choice is more important than technology choice

Another clear dividing line is how operators choose partners.

Here's the problem with many AI suppliers in transport right now:

they're selling advanced capabilities without understanding how transport actually runs day-to-day. The promises sound transformative. The reality? Tools that don't fit operational workflows, implementations that stall, and adoption that never quite happens.

The operators moving fastest work with partners who've actually run transport operations. Instead of buying technology in isolation, they're solving specific operational problems. Scheduling that reflects actual demand. Maintenance based on vehicle condition. Disruption management that happens in minutes, rather than hours. And customer contact that reduces repetitive queries rather than just handling them faster.



One tip we often give to clients is a simple test to separate serious partners from technology vendors:

Does the conversation start with understanding how your operation runs or with a product demo?

If a partner leads with features before they understand your constraints, your routes, your people, and your pressures, they're selling technology. If they start by asking how scheduling actually works in your depot on a Tuesday morning in January, they might actually be able to help.

Where leaders are focusing first

For COOs and CIOs, progress starts with clarity on the operational problems that need attention. From there, specific use cases emerge where better use of existing data can support improved outcomes.

Leaders who make progress typically:

- Agree on a small number of priority operational issues
- Put governance in place early to manage risk and accountability
- Use pilots to demonstrate return and build confidence
- Address skills and capability alongside deployment

The aim is to build momentum and confidence through delivered improvements, rather than waiting for everything to be fully designed or agreed.

Moving past the infrastructure vs AI debate

Now, many organisations still frame investment decisions as a choice between fixing ageing infrastructure and deploying new digital capability. **In practice, this is a false dilemma.**

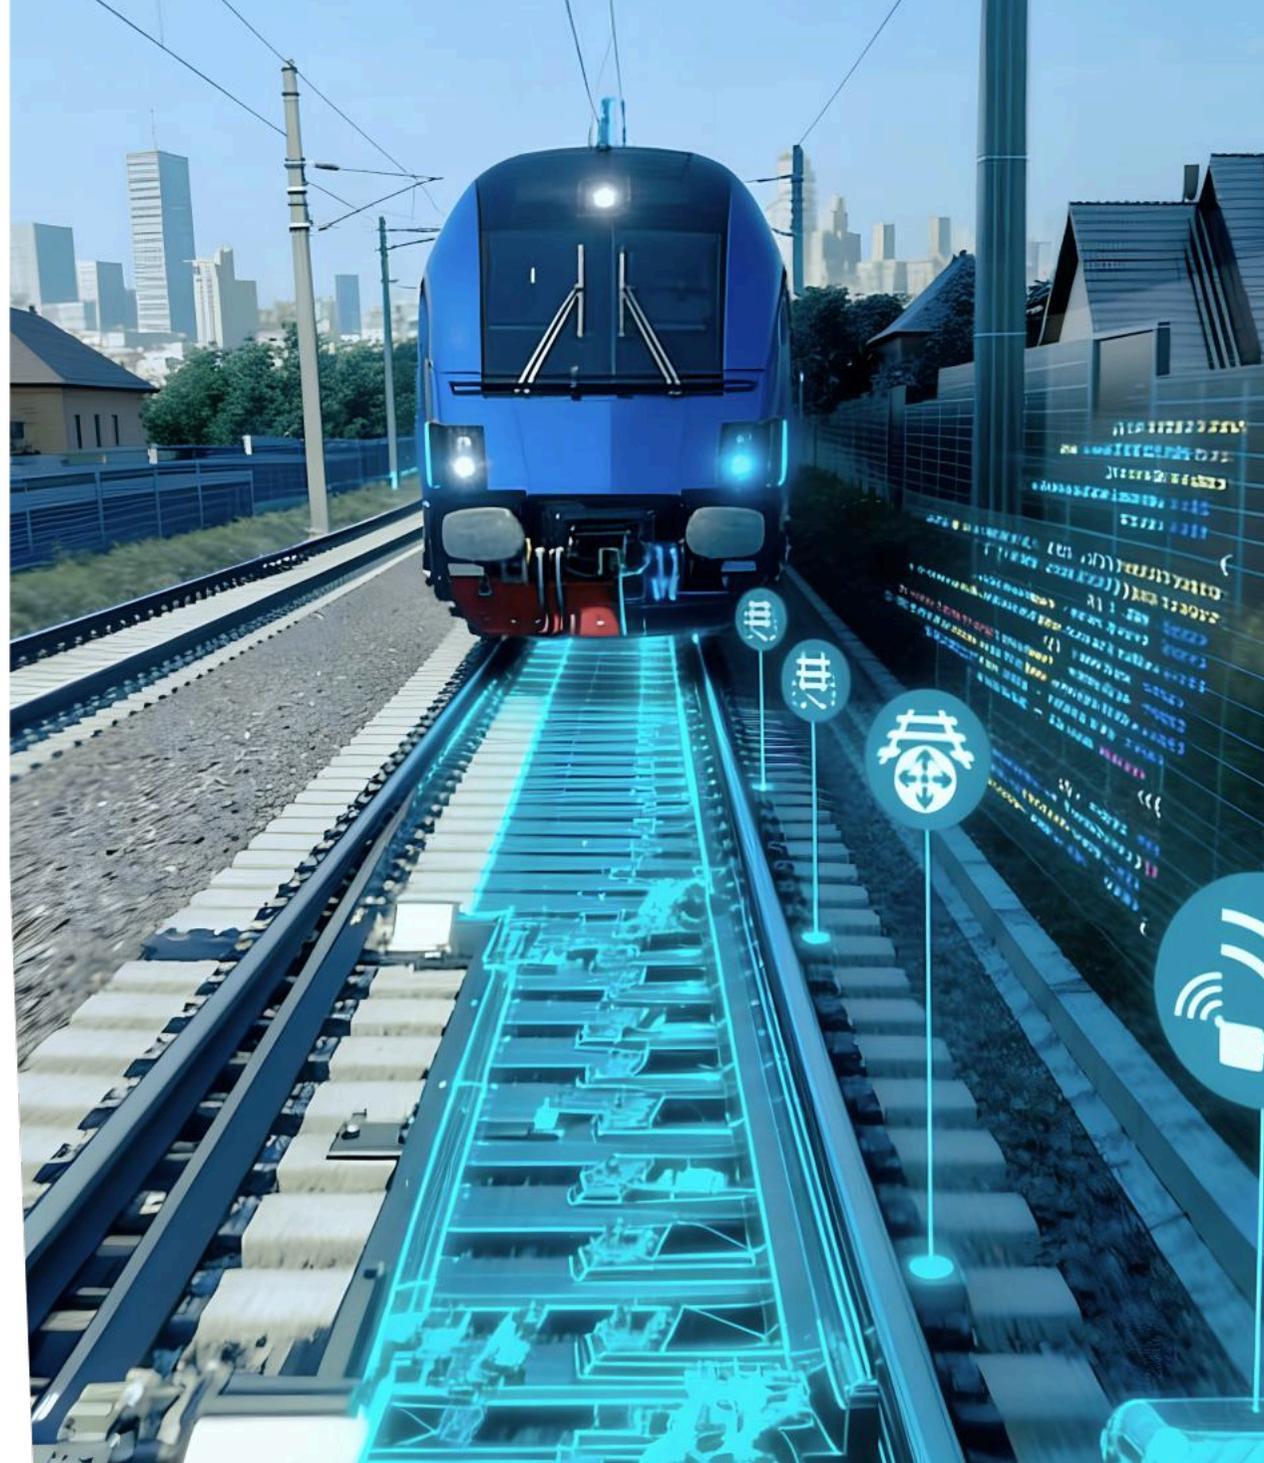
Waiting for perfect infrastructure often becomes a way of delaying change. (Perfection rarely arrives.)

The operators who are making progress fix what is genuinely critical, such as safety, compliance, and major failure risks. And at the same time, they invest in modern integration layers that allow legacy systems and newer tools to coexist. Then, they deploy AI against use cases that can deliver a return within 12 months, using the systems they already have.

Leaders often describe this balance as:

- Around **30%** on foundation stability
- Around **20%** on integration
- Around **50%** on AI use cases that deliver near-term operational return

This balance allows organisations to improve resilience while continuing to move forward.



Let's look at another practical example...

Real AI use case:

Predictive maintenance for bus fleet reliability

Public Transport Authorities increasingly expect **demonstrable reliability**, rather than reactive explanations after failures occur. For many bus operators, traditional telematics systems are no longer sufficient.

While a major bus operator's existing technology provided solid operational control, it lacked the learning capability needed to **predict failures before they occur**. This gap is becoming more visible as **zero-emission fleets introduce new and less predictable failure modes**, and as contract performance is measured through **availability, cancellations, and lost mileage**.

What was implemented

An **AI-powered predictive maintenance platform** was introduced as a complementary intelligence layer alongside existing telematics.

The platform uses applied models trained on **real-world bus fleet behaviour** to detect developing faults **days or weeks earlier** than driver reports or workshop diagnostics. Hardware integration provides the data depth required across **ICE, hybrid, and zero-emission vehicles**.

Operational impact

Workshops can prioritise critical interventions, align maintenance activity with actual vehicle condition, and improve fleet availability.

Maintenance shifts from reactive recovery toward **evidence-based reliability management**, changing performance discussions with PTAs from explanations to outcomes.

The results?

Benefits seen to date include:

- **Over 40% reduction in vehicle recoveries**
- **An average of 1 critical issue identified proactively per bus, per month**
- **Lost mileage reduced by almost 1.5%**
- **Labour time per bus reduced by almost 70%**

As franchise models give authorities greater responsibility for bus assets, the ability to demonstrate early fault detection and proactive maintenance is becoming increasingly important to contract performance.

Leading transport operations in 2026

In 2026, pressure across transport operations doesn't let up. Disruption, congestion, rising costs, driver shortages, safety performance, ageing assets, and technology debt are all competing for the same limited time, budget, and leadership attention.

The operators handling this best have changed when and where decisions happen. Action happens earlier. Coordination improves. Decisions move closer to where problems actually emerge, with data and automation supporting experienced teams to act sooner and more consistently.

Crucially, they're not chasing moonshots. Maintenance goes from scheduled to condition-led. Capacity adjusts closer to actual demand. Customer contact focuses on resolution, not endless repetition. And safety risks get flagged and addressed earlier in the cycle, before they become incidents.

At the same time, leadership priorities are clear. Critical infrastructure is stabilised. Operational improvement that can be delivered now is pursued using existing systems. Change is managed in parallel, not held back by sequential programmes.

Partner choice reflects this approach. Operators are working with organisations that understand how transport runs day to day, applying new capabilities in ways that fit existing operations.

This is how leading transport operations are being run in 2026.

Reviewing your operation under pressure?

[Mind The AI Gap](#) works with transport operators to apply AI to real operational problems.

We're led by people who've run transport operations, not just consulted on them, including senior operational roles and advisory work with bus and rail operators. That experience shapes how we approach planning, maintenance, customer contact, and disruption management in practice.

We're here to work alongside your experienced teams, using the systems and data you already have, to help you make better decisions earlier and get more consistent outcomes.



**MIND THE
AI GAP**

If you're reviewing your operational strategy for 2026 and want a conversation that starts with how your operation actually runs, **let's talk:**

 info@mindtheaigap.com

