





optibus

AGENTIC AI FOR PUBLIC TRANSPORTATION

A practical guide to today's landscape,
real challenges, and what works now

What's inside

- 
- 1 Introduction**
Why solid planning still turns into slow execution, & what's changing.
 - 2 An assistant answers. An agent acts.**
The distinction that matters, & the technical shift that made it possible.
 - 3 Why transit is uniquely hard for AI to get right**
Massive datasets, human unpredictability, & why generic tools underdeliver.
 - 4 Where agentic AI adds value across the operational cycle**
Planning, scheduling, operations, & control — with real-world proof points.
 - 6 The questions every agency should ask**
Jobs, data safety, & trust — addressed directly.
 - 7 Where this is headed**
From discrete tasks to a coordinated team of specialist agents.
 - 9 How Optibus is building toward this**
Optibus Agent, embedded across planning, scheduling, allocation, & control.
 - 10 Conclusion**
- 

The Gap Between Planning and Executing

Public transportation runs on decisions. Thousands of them, every single day: which route gets an extra trip, which driver covers an open shift, which detour keeps a line moving during a road closure, which schedule change actually holds up once real traffic and real people are added to it.

The tools to make good decisions have existed for years. Scheduling software, GTFS feeds, real-time vehicle tracking, optimization engines: the industry has invested heavily in systems that model complexity well. And yet, agencies everywhere report the same experience: the planning is solid, but the execution is slow. Pulling the right view takes longer than it should. Cross-referencing driver availability against work rules eats a whole shift. A disruption that should take minutes to resolve instead takes an hour, because the information needed to make the call is scattered across five screens and two spreadsheets.

That gap, between what agencies can plan and what they can actually execute in the moment, is where a new generation of artificial intelligence is starting to make a real difference. Not the general-purpose chatbots most people have already tried, but a more capable, more specific category: agentic AI, built to work inside the operational reality of public transportation.

This ebook explores what that actually means: a grounded explanation of the difference between AI that answers and AI that acts, where that distinction matters most in transit operations, what the honest limitations are, and where the industry is realistically headed over the next few years.



An assistant answers. An agent acts.

Most people's mental model of AI comes from tools like ChatGPT, Claude, or Gemini: you type a question, you get an answer. That's genuinely useful, but it's also a fairly narrow slice of what's now possible, and it's important to be precise about the difference, because the terms "AI" and "agentic AI" get used almost interchangeably in vendor marketing, and they shouldn't be.

An assistant answers a question

You still do the work. Ask it to summarize a report or explain a rule, and it will, but the follow-through, the clicks, the actual changes in a live system, are still on you.

An agent works autonomously toward a goal

It uses tools, data, and system access, across multiple steps, until the job is actually done. It doesn't just tell you which driver is available; it can check the roster, cross-reference the rules, and take the action, with a human confirming before anything goes live.

How an agent closes a request, start to finish



That distinction is the line that matters for public transportation. A chatbot that can describe how scheduling works is a nice-to-have. A system that can actually go into a live schedule, apply a rule change, and show you the resulting trade-offs is a fundamentally different tool, because it closes the distance between deciding and doing.

It's also worth being honest about why this wasn't possible until recently. Classic optimization software needs clean, structured input to work well. Feed it messy, incomplete, or inconsistently formatted data, and it either fails outright or produces something unreliable. Large language models flipped that constraint on its head: they're specifically good at making sense of unstructured, messy, real-world data, including mismatched formats, incomplete records, and plain-language descriptions of intent, and turning it into something a structured system can act on. That's the technical shift that made agentic AI in transit possible in the first place.

Why transit is uniquely hard for AI **to get right**

Running a transit operation is one of the more complex logistics problems that exists, precisely because it sits at the intersection of massive datasets and constant human unpredictability.

Consider what a single scheduling decision actually depends on:



What's happening on the road right now



Labor agreements with dozens of interacting rules



Every change ripples into a dozen others



Where people live & work



Traffic patterns



Driver availability & qualifications



Every vehicle's condition & assignment

Thousands of decisions like this get made daily, and until recently, nearly all of these decisions rested on the judgment of experienced planners, schedulers, and controllers. Software helped, but it didn't help enough.

Specificity, more than any raw model capability, is what separates a system that's genuinely useful in daily operations from one that looks impressive in a demo and stalls in production.



This is exactly why generic, horizontal AI tools tend to underdeliver in transit specifically. A general-purpose assistant has no native understanding of what a "split shift" is, why an Hours-of-Service violation matters, or how a labor agreement's break rules interact with a driver's seniority. It can be prompted and configured toward those concepts, but it starts from zero, and getting it to a useful level of domain fluency takes months of integration work that most agencies don't have the bandwidth to do.

Purpose-built agentic AI skips that gap. It arrives already fluent in the vocabulary and constraints of transit operations, including what a duty is, what a GTFS feed contains, and what "coverage" means for a route, so it can reason about trade-offs instead of merely retrieving information.

Where agentic AI adds value across the operational cycle

Agentic AI isn't a single feature. It's most useful when it's woven across the full lifecycle of running a transit network. Four areas stand out as where the technology is delivering real value today.

One operational cycle, four hand-offs

PLANNING > **SCHEDULING** > **OPERATIONS** > **CONTROL**


Planning

When a stop goes inactive or a road closes, planners traditionally have to manually trace every affected route and pattern, then fix each one individually. An agent that understands the network structure can identify every affected pattern at once, propose replacements, and apply changes across the board. Riverside Transit Agency in California has already piloted a version of this: a disruption management tool that automatically pushes real-time detour updates to riders and operators, cutting response times and improving schedule reliability.

Scheduling

Before a scheduler can do any real scheduling work, they usually have to get oriented first, finding the right view, filtering to the right routes, and pulling the right statistics. An agent that understands plain-language requests can filter a schedule instantly, or answer questions about duty stats, vehicle stats, and coverage figures directly, in context, without an export-and-search detour.

Where agencies plan to deploy AI next — 2026 APTA member survey



Share of 32 surveyed agencies indicating current or planned AI use, following customer support and analytics as the areas of heaviest current deployment.

Operations

Operations teams juggle three parallel burdens every day: knowing what already happened (audit trails), knowing what's currently wrong (issue flags), and filling the gaps in coverage (driver assignment). Each is its own manual process today. Predictive maintenance is one of the clearer proof points of what AI can already do here, and agentic AI extends the same logic to plain-language questions, such as "who approved this absence, and when," or surfacing the right available driver by cross-referencing qualifications, history, and rest requirements, turning what used to take ten minutes of digging into a ten-second answer.

Control

Nowhere does time matter more than in real-time control. A dispatcher needs to know which vehicles are running late, locate a specific bus using whatever identifying detail they happen to remember, and get a message to a driver, all under pressure, often during the worst part of the day. An agent that can chain these actions together (locate the most delayed vehicle and message the driver, in one exchange) removes the friction of switching between screens and systems exactly when speed matters most.

75%

increase in maintenance productivity after New York's MTA applied AI-powered predictive maintenance to its bus fleet — with material costs down 24 percent.

22 to 787

bus lane violation citations over a comparable two-month period after AC Transit implemented AI image recognition for enforcement.

Across all four areas, the pattern is the same: the agent doesn't just hand over more information for a person to interpret. It prepares a decision, with the reasoning and trade-offs already laid out, so the human's job shifts from gathering to judging.

This isn't a single-vendor trend. The direction of travel across the industry is consistent, even as individual agencies are still at different points along it.

The questions every agency should ask, and the honest answers

Any serious conversation about AI in public transportation has to address three concerns directly, because agencies are right to raise them.



"Will this replace my team?"

No, and this deserves a clear-eyed answer, not a reassuring platitude. Running transit well requires human empathy, local knowledge, and judgment calls that don't reduce to a formula: how to balance rider needs against operating costs, how to read a community's tolerance for service changes, how to weigh a trade-off that a system can describe but shouldn't decide.

What agentic AI removes is the administrative burden around those decisions, the manual searching, cross-referencing, and data-wrangling that eats the hours a planner would rather spend on strategy. A junior planner gains access to the reasoning of a much more experienced one; a senior planner gets freed from repetitive correction work to focus on network design. Nobody's job disappears; the work each person does shifts upward.



"Is our data safe?"

This is the single most important operational question, agencies should push for specific answers, not vague assurances. The responsible standard: data stays isolated within a closed, secure environment; it is never used to train public or shared models; and any AI processing happens without customer data leaving the platform it originated in. If a vendor can't answer precisely how data flows, where it's processed, and whether it's retained, that's a signal to keep asking.




"Can I trust the output?"

This is where the technology needs to be paired correctly. Pure generative AI, on its own, can produce confident-sounding but incorrect output, a real risk when the subject is a live transit schedule. The responsible approach couples generative reasoning with deterministic, rules-based optimization engines that have been validated for years: the AI handles translating intent and messy data into the right parameters, while a separate, proven engine handles the actual mathematical optimization. And critically, a human stays in the loop: the agent proposes, explains its reasoning, and a person decides before anything is applied to a live system.

Where this is headed

What exists today is a meaningful but still early stage of a much larger shift. It's worth being honest about that distinction, because overpromising erodes exactly the trust this technology needs to earn.

Today, agentic AI in transit performs discrete, high-value tasks: filtering a schedule, resolving a data import error, finding an available driver, locating a vehicle. Each is useful on its own. But the more interesting trajectory is what happens as these capabilities start working together.



★ From tools to a team of experts

Instead of operating a series of separate systems, transit teams start directing a coordinated set of specialist agents that hand off to one another across the full lifecycle of a service, from initial network design through to a bus on the road in real time. The person sets the intent and the standards; the agents handle the coordination between them.

★ A network that improves itself

Rather than reacting to a disruption after it happens, an operation that continuously reads every available data source can start anticipating problems, flagging what's about to go wrong before it does, proposing a fix along with the reasoning behind it, and learning a team's preferences over time. The "data-gathering tax" that currently consumes so much of a planner's day gradually disappears, leaving more time for judgment and design.

★ An expert embedded in a wider ecosystem

As AI tooling becomes more interconnected across an agency's broader technology stack, a transit-specific agent doesn't operate in isolation. It becomes the domain expert that the rest of an organization's systems can call on, the piece that actually understands how public transportation works, plugged into a larger network of tools.

This trajectory isn't unique to transit software vendors, either. Transportation regulators and public agencies are moving the same direction:

The U.S. Department of Transportation has signaled a strategic shift toward agentic AI for tasks like grant-compliance review and analyzing complex weather and traffic datasets to generate real-time alerts, the same underlying logic of an agent that prepares a decision rather than just retrieving information, applied one level up, at the policy and oversight layer.

None of this replaces the fundamental purpose of the work. It's still about getting people where they need to go, reliably and affordably. The technology is a means to serve that goal better, not a goal in itself. If it succeeds, the measure won't be how sophisticated the AI looks in a demo; it will be whether transit becomes more reliable, more efficient, and a genuinely better option for the people who depend on it.



How **Optibus** is building toward this

Everything above reflects where the broader industry is heading, a shift already underway across multiple vendors and platforms, each approaching it from different angles. At Optibus, this is the direction we've built toward directly: **Optibus Agent**, an AI agent purpose-built for public transportation.

Optibus Agent is embedded across every stage of the Optibus platform rather than bolted on as a separate chat window:

**Planning | Scheduling | Driver Allocation | Compliance Monitoring |
Control Room Functions | Dispatch | Operational Reporting**



In practice, that means:



In planning and scheduling, the agent helps build duties that align with business objectives, collective agreements, and driver preferences, for example generating balanced rosters that avoid overtime or split shifts where possible, and helps planners evaluate route and timetable options against passenger demand.



In driver allocation, it can identify qualified staff to cover an open shift or a last-minute absence by cross-referencing availability, work history, and qualifications, the kind of task that used to mean digging through several systems by hand.



Preference Designer lets schedulers type a rule in plain language, such as "no more than ten duties over nine hours," and turns it into accurate, ready-to-use scheduling logic, work that traditionally required configuration expertise or custom coding.



Schedule Analysis compares optimized scenarios against each other and produces presentation-ready trade-off summaries, so planners can bring stakeholders a clear picture instead of a raw data export.



In control and dispatch, it supports real-time operational monitoring and reporting, so controllers can interrogate what's happening on the network in plain language instead of hunting across dashboards.



Two things are consistent with the principles laid out earlier in this ebook. First, the agent doesn't replace the underlying optimization: it sits on top of the same deterministic, rules-based scheduling engine that's powered Optibus for over a decade, so the reasoning stays fast without sacrificing the mathematical rigor transit scheduling requires. Second, a human stays in the loop at every step: the agent proposes and explains; a person decides before anything goes live.

CONCLUSION

Closing the Gap

Public transportation has never lacked data or planning tools. What it has lacked is a way to turn them into timely action, without the manual gathering and cross-referencing that slows every decision down. Agentic AI is the first technology built specifically to close that gap: it prepares decisions, explains its reasoning, and keeps people in charge of what goes live.

The agencies and operators that gain the most will be the ones that start now: ask vendors the hard questions about data security and human oversight, and choose tools that already speak the language of public transportation.

Ready to see what an agent built for public transportation looks like in practice?

EXPLORE OPTIBUS AGENT

About Optibus

Optibus provides a unified software platform for public transportation planning, scheduling, rostering, operations, and passenger information. Agencies and operators in thousands of cities worldwide use Optibus to harness the power of artificial intelligence, optimization algorithms, and cloud computing and offer more efficient, sustainable, and equitable public transportation services.

www.optibus.com | info@optibus.com