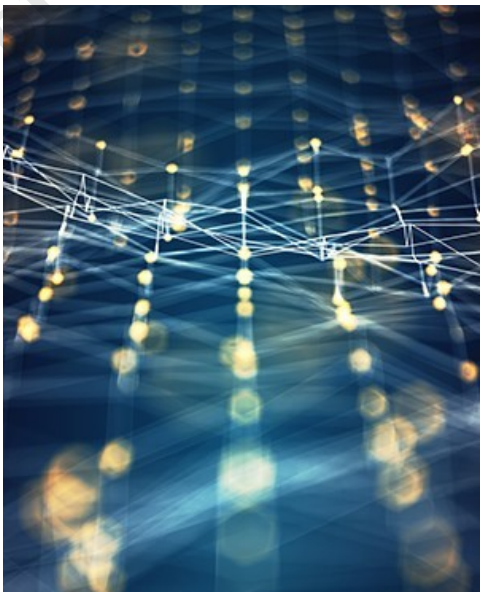


www.pipelinepub.com
Volume 22, Issue 12

The Future of Traffic Isn't Robots. It's the Network.

By: Timothy Menard



In August, a Pennsylvania state lawmaker began drafting a bill that would require a licensed human driver aboard any automated transit vehicle weighing more than 10,000 pounds, even though the transit agency the legislation is aimed at says the technology is nowhere close to arriving in the city. The proposal would not ban automated buses outright. It would simply keep a human in the seat until trust catches up with the technology.

That single bill says more about the current moment in transportation than any product launch could. The public conversation about AI and traffic has become almost entirely a conversation about robots that drive. It is not the conversation about the AI that is already driving results.

A Crowded, Noisy Narrative

The unease is not irrational. In Gothenburg, Sweden, a new autonomous bus service was rear-ended by a tram during its second passenger trip, an incident that made international headlines despite causing no serious injuries. In Austin, one automaker has disclosed more than a dozen crashes involving its robotaxi fleet since the service launched in 2025, and another operator recalled over a thousand vehicles after software caused collisions with gates and chains. Each incident is minor on its own. Together, they are shaping a public narrative that AI behind the wheel, whether in a car or a bus, is reckless or unready.

That narrative is a reasonable response to a genuinely hard problem: teaching software to make split-second physical driving decisions with no fallback. But it is a narrative about one specific application,

Not for distribution or reproduction.

© Pipeline Publishing, L.L.C. All Rights Reserved.

not about artificial intelligence and transportation as a whole. Meanwhile, congestion keeps getting worse almost everywhere. Chicago now leads the nation in time lost to traffic, with drivers there losing 112 hours a year, and nearly nine in ten U.S. cities saw delays increase over the previous year. Traffic is not waiting for the robotaxi question to be settled.

Two Different Jobs, One Confusing Label

It helps to separate what is actually being asked of "AI" in these two cases. An autonomous vehicle is AI making continuous, real-time physical decisions, replacing a human driver entirely. It is one of the hardest open problems in engineering, and it is still, visibly, a work in progress. AI-powered adaptive traffic signal control, by contrast, uses artificial intelligence for something far less dramatic: adjusting how long a light stays green based on real conditions, and increasingly, predicting what traffic will look like several intersections ahead.

One is trying to replace the driver. The other is trying to replace a stopwatch. Not every signal timing system delivers as promised, and poorly integrated sensor data can leave even a smart intersection performing like a dumb one. But the narrow, already proven uses of AI in traffic management should not be judged by the struggles of a much harder and unrelated problem.

The Quiet Version That Already Works

This is where the less exciting story deserves more attention. AI-powered adaptive signals have cut travel times in many regions and reduced emissions by cutting the idling that comes with unnecessary stops. Federal research simulating an intersection in Charlotte, North Carolina found that AI-based transit signal priority reduced average delay per person by nearly 19% during peak hours, a measured outcome rather than a projection.

None of this requires a vehicle that drives itself. It requires software that gets a little better at anticipating what is about to happen at an intersection and nudging a signal a few seconds ahead of time. It is the traffic equivalent of a thermostat that learns a household's schedule, not a robot that walks the dog. And it does not require new fleets, new safety regulations, or years of public trust building. Cities already own the signals, and most of the infrastructure consideration is already won.

The Edge Hardware Question

There is a second, less visible debate underneath the first, and it concerns where the intelligence should physically live. For years, the default assumption in smart intersection projects was that more capability meant more equipment at the curb. Processing units, dedicated radios, and extra cabinets were added at each location so decisions could be made close to the street. That approach works, but it carries real costs. Every box added to a pole or cabinet has to be purchased, installed, powered, maintained, and eventually replaced, and a city with hundreds of intersections multiplies each of those tasks hundreds of times.

Heavy edge hardware also tends to age quickly, because the software running on it improves faster than the devices can be swapped out. It creates a patchwork too. A corridor that crosses two jurisdictions may end up with different equipment generations and update schedules, which makes coordinated timing harder rather than easier.

A Lighter Architecture

A growing alternative treats the existing roadside controller as the stable foundation and moves the heavy computation off the street. Instead of replacing the controller, the software connects to it, reads the signal state it already produces, and sends back timing adjustments. The analysis of traffic patterns, the prediction of arrivals, and the coordination between intersections run in a cloud

Not for distribution or reproduction.

© Pipeline Publishing, L.L.C. All Rights Reserved.

environment, where processing power can be scaled up without touching a single cabinet.

In practice, this keeps the physical footprint small. The controller continues to handle the safety-critical work it was certified to do, including conflict monitoring and the guarantee that opposing directions never receive a green light at the same time. The cloud layer supplies the intelligence that sits above that foundation. Because improvements arrive as software updates, a better prediction model reaches every intersection at once.

Why Edge Networking Decides the Outcome

None of this works without a dependable network, and that is where the real engineering challenge lies. Edge networking is the connective tissue between vehicles, signals, and the cloud. A bus that is two blocks away needs its signal request recognized in time to extend a green, and an ambulance approaching an intersection needs clearance before it arrives, not after. Both depend on latency, the delay between an event on the street and a response from the system.

This is where vehicle-to-everything data, commonly called V2X, becomes valuable. Vehicles and infrastructure can share position, speed, and priority status directly, giving the system an earlier and more precise picture than a pavement sensor alone. When that data feed is combined with the controller's own signal information and routed over a well-designed network, the delay between detection and action can be kept low enough for emergency preemption and transit signal priority to work reliably. A city pilot in Winnipeg estimated that preemption cut fire and ambulance response times by 15 to 20 percent, though the report noted that roadway construction prevented a direct comparison.

The point worth emphasizing is that low latency does not have to mean heavy hardware at every corner. Careful network design, efficient data handling, and a clear division of labor between the controller and the cloud can deliver fast responses without turning each intersection into a small data center. Resilience matters as well. If a connection drops, the controller must keep running safely on its own and fall back to standard timing plans.

What Cities Gain

The practical benefits of this approach show up in places municipal leaders care about. Deployment costs fall when fewer devices need to be bought and installed. Rollouts move faster because software needs little roadwork. Corridors that cross city lines can be coordinated under a single logic, so a green wave does not break at a municipal boundary. And because the same architecture handles routine traffic, emergency vehicles, and transit, throughput improves across the whole network rather than at isolated showcase intersections.

Cities should still ask about data ownership, cybersecurity, and uptime, and test interoperability with existing controllers rather than assume it.

What Is Actually at Stake

None of this is an argument against caution on autonomous vehicles. The safety questions raised by recent incidents are legitimate and deserve serious scrutiny. But public trust in "AI and traffic" is currently being shaped by the hardest, highest-risk version of the technology, at precisely the moment cities need to make calm, evidence-based decisions about the easiest, lowest-risk version. Congestion and emissions are not pausing while that debate continues. Every year a city delays adopting smarter signal software is a year of avoidable delay and fuel burned that a decades-old traffic light, running better software, could have prevented.

Not for distribution or reproduction

© Pipeline Publishing, L.L.C. All Rights Reserved.

The future of traffic that actually arrives on time will not look like science fiction. It will not get rear-ended by a tram on its second day in service. It will be a traffic light that already knew a driver was coming and turned green when it needed to. No one will post a video of that. It will just work.

Not for distribution or reproduction.

© Pipeline Publishing, L.L.C. All Rights Reserved.