NOCoE
National Operations Center of Excellence

Peer Exchange Report

AI for TSMO

Sacramento, California
May 19–20, 2026

Executive Summary

In May 2026, the National Operations Center of Excellence (NOCoE) brought 20 transportation agencies and national partner organizations to Sacramento for two days of candid, working-level conversation about the use of artificial intelligence (AI) in transportation systems management and operations (TSMO). Hosted by Caltrans, the exchange was the in-person follow-up to a 2025 virtual peer exchange and it marked a clear turn in the national conversation: from *whether* and *why* agencies should use AI to *how* they stand it up, pay for it, contract for it, govern it, and staff it.

Five key learnings stood out across the two days:

- **Data readiness, not model choice, determines success.** The agencies furthest along: Caltrans, with two production AI platforms live within two years of a gubernatorial executive order, and Texas DOT, whose multi-agent data tool projects millions of dollars in annual staff-time savings, both attribute their progress to years of unglamorous data standardization and governance, not to any particular AI product.
- **The field is organizing itself from the bottom up.** Multi-state pooled funds, a new State DOT Council for Strategic AI Adoption, a federal AI Center of Excellence, and new association resources are converging into the shared clearinghouse agencies say they need.
- **Governance is converging on clear principles: decision support, not automated.** Agencies are keeping a knowledgeable human responsible for every consequential output, writing that accountability into policy and contracts, and insisting that expert review be genuine rather than a rubber stamp.
- **The leading constraint has shifted from technology to funding.** Participants named funding, specifically, the absence of dedicated, sustained money for work that sits outside the traditional capital program, as the primary constraint.
- **Success isn't the only valuable tool to progress.** Iowa DOT presented pilots that did not work alongside those that did, and the room treated those post-mortems as among the most useful content of the exchange.

The body of this report walks through each session of the Peer Exchange in the order of the agenda, with figures drawn from the presenters' own slides. It closes with consolidated key learnings, the next steps to advance the adoption of AI for TSMO, and the public resources referenced during the event.

Table of Contents

Background

The Peer Exchange was the in-person follow-up to the virtual AI for TSMO peer exchange NOCoE convened on June 24, 2025. That earlier Peer Exchange drew 43 practitioners from 28 agencies and established a baseline understanding of the field. Interest was strongest in traffic management and optimization and in data analysis and decision-making, but agency maturity varied enormously. Some agencies were already piloting tools in production; others were still in strategic ideation. A 2025 poll of participating agencies found that 14 of 29 had an AI strategy in development, only two had a fully developed strategy, and twelve had no published AI policy at all. The closing sentiment in 2025 was an invitation to "fail fast and learn quickly," and a clear demand emerged for a deeper, hands-on conversation about how agencies stand up, scale, govern, contract for, and staff AI.

Where the 2025 session was largely about awareness, the 2026 exchange was about implementation in practice. The agenda was built to move the group from the descriptive to the operational. Day 1 established what agencies are doing, including the host's vision, Caltrans' production tools, multi-state coordination models, the resources offered by national partners, and an all-agency roundtable. Day 2 turned to the harder questions of how to scale and sustain that work, with sessions on expanding from pilot to program, contracting for AI, and on building the workforce and the training behind it. Each session paired short presentations with extended discussion, and two roundtables gave every agency in the room time to speak.

Day 1 — Tuesday, May 19

Welcome and Purpose

Nicholas Ramfos of NOCoE opened the Peer Exchange by connecting the event to the conversation begun a year earlier. He reminded participants of the June 2025 virtual session and proposed a definition of what success would look like. Rather than a comprehensive survey of every tool on the market, he asked each participant to leave with two or three concrete ideas worth pursuing at their own organization. He walked briefly through the 2025 findings: the maturity spread, the small number of agencies with finished strategies, and the recurring 2025 concerns of cost, organizational readiness and risk aversion, data quality and privacy, human oversight, and accountability.

Lastly, he noted that the conversation was meant to continue beyond the two days in Sacramento, and that NOCoE wanted to hear what form of continued support would be most useful, whether webinars, additional peer exchanges, a best-practice compilation, or a shared library of use cases.

Host Welcome and the Caltrans AI Vision

Introduction to the Conversation

Kevin Riley of Caltrans welcomed the group as host. He framed California's position with candor: Caltrans does some elements of TSMO very effectively, but the agency had not yet looked at the full scope of what AI could make possible. The ability of AI to synthesize and absorb data, he suggested, would surface insights at a speed and scale not previously achievable. He illustrated this point with the observation that technology could now produce in minutes analyses that once took staff weeks of effort, sometimes revealing patterns in the system that even experienced District leaders had not seen. He then handed the floor to Vidhu Shekhar, who leads Caltrans' data and AI program, to describe how the agency had organized itself to pursue that potential.

The Caltrans Data and AI Program

Shekhar's central message was that data, and AI must be treated as a single undertaking. Data, he argued, is the lowest common denominator beneath everything an agency does with technology; AI initiatives that are not built on a sound data foundation will not succeed. For that reason, Caltrans had deliberately framed its work around use cases rather than around models or vendors; the question the program asks first is what problem needs solving, not which tool to buy. In standing up the program, Caltrans looked to peer states such as Texas and Utah as models for policy, goals, and structure.

The program's origin was an executive order at the state level, with sponsorship cascading from the agency Secretary through the Director to the program leaders. California's position as the epicenter of the AI industry includes a large share of the country's leading AI companies located within a short distance of Sacramento which gave the effort a particular urgency. Shekhar described his own role as that of an internal consultant and, in his words, the chief cheerleader for the tens of thousands of Caltrans employees across every program area, from highways to aeronautics to permitting and sustainability. He returned often to a metaphor of "islands of innovation" scattered across a large organization, and to the work of building bridges between them; governance, in his framing, exists to keep people from falling off the edges, not to fence them in.

What distinguished the Caltrans account was the institutional machinery the agency built around the work, not simply the tools it produced. The Caltrans program traces its authority to a directive from the California State Transportation Agency (CAlSTA) establishing a **Chief Data and AI Officer (CDAO)**. The CDAO's mandate is to drive integrated data governance and AI adoption, to ensure that implementation is ethical, compliant, and transparent, to align AI initiatives with organizational goals across departmental lines, and to define the metrics by which the impact of those initiatives is measured. A Data and AI Strategy Office provides enterprise-wide leadership for data and AI policy, standards, and practices. For a TSMO audience, it is important to note that this responsibility spans both information technology and operational technology rather than stopping at the IT boundary.

The program organizes its work into five workstreams:

- the design of the CDAO role and organization;
- strategy and roadmap development;
- governance and guidance;
- use cases,
- technology, and tools;
- and communications and organizational change management.

Caltrans built governance itself as a three-tiered model. A Director's Policy establishes the CDAO's authority over AI policy, governance, and strategy. A Deputy Directive sets the operational requirements for acceptable use, satisfying the state's administrative-manual requirements for AI. Beneath those, domain-specific standard practices translate the directive-level requirements into operational procedures. The agency's strategy is built on a deliberately short two-year horizon and is paired with an objectives-and-key-results framework, an approach Shekhar acknowledged is unusual for a public-sector strategic plan.

Caltrans was equally candid about where it stands. A readiness assessment published in October 2025 scored the agency against a five-level maturity scale across five dimensions: strategy, people, process, technology, and data. The overall score was 1.92 out of 5, placing Caltrans in the "tactical" band, with individual dimensions ranging from a low of 1.58 for process to a high of 2.21 for both people and technology. The agency stood up a generative-AI subcommittee in September 2025 to begin closing those gaps, and named its top maturity opportunities as governance execution, data strategy and access, portfolio discipline, continuity and cost management, and workforce enablement.

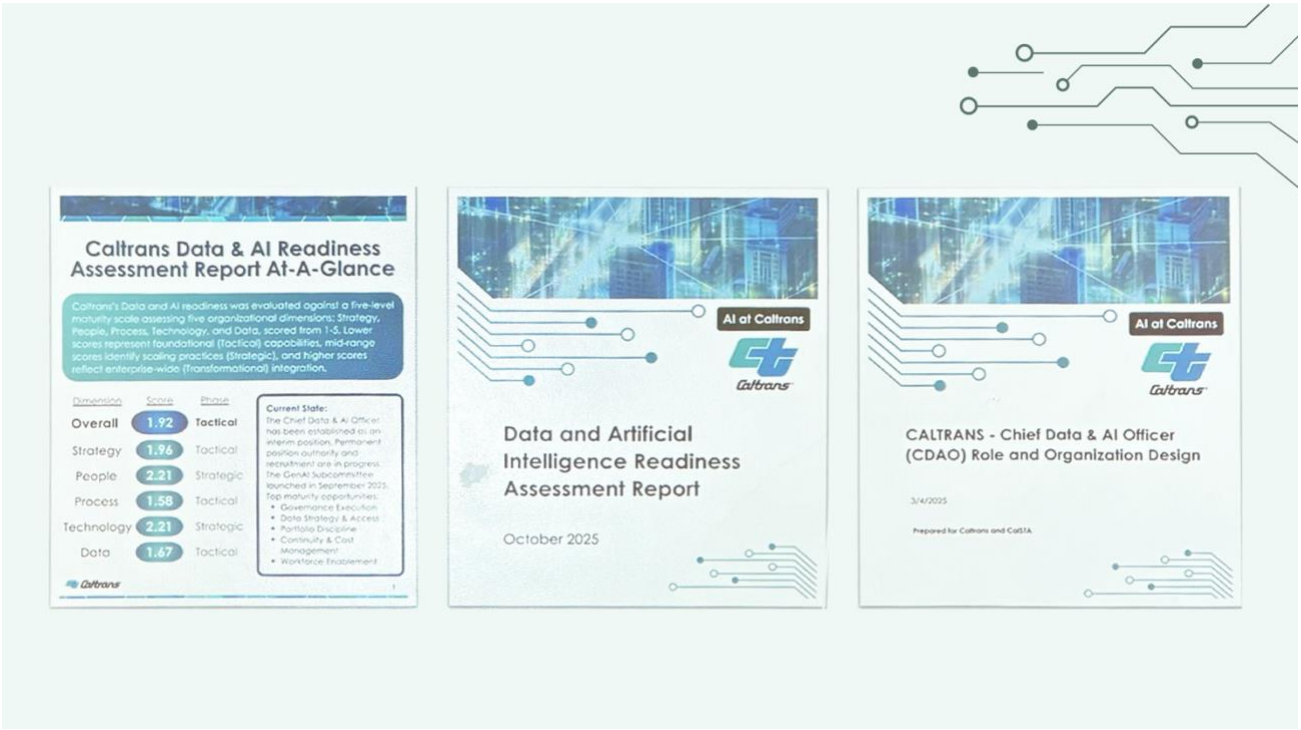

Figure 1. Caltrans' published self-assessment of its data and AI readiness (October 2025): an overall 1.92 on a five-point maturity scale, with dimension scores for strategy, people, process, technology, and data. The candor of the baseline shaped the governance buildout described in this section. Source: Caltrans presentation.

To manage the flow of ideas into that governance structure, Caltrans developed an AI Use Case Framework, a four-step pipeline that moves a candidate concept from intake and pre-qualification through prioritization to approval, with three entry points for new project requests, integrations into existing solutions, and individual staff ideas. A 2026–2028 Data and AI Strategic Plan captures the agency's direction, organized around five goals:

- improving safety,
- mobility, and business outcomes;
- empowering a people-first workforce;
- building a trusted data foundation;
- delivering measurable efficiency and cost savings;
- and governing AI responsibly and sustainably.

Among the program's nearer-term ideas was a notion Shekhar called data minimization: the recognition that an AI system often needs the *right* curated data rather than the largest possible volume of it. The inventory of concepts moving through the use-case framework is substantial as more than 300 candidate use cases were developed, advancing through a pipeline that runs from discovery to proof-of-concept testing, piloting, and development. The active portfolio already reaches well beyond the two flagship traffic tools, taking in an internal generative-AI chatbot on the agency's intranet, the enterprise productivity assistant, adaptive AI traffic-signal control, AI-enabled wrong-way-driving detection, an AI-enabled learning-management system, and work-zone safety applications.

The discussion that followed previewed many of the event's recurring tensions. Asked who owns **cybersecurity** for these systems, whether responsibility begins in IT, Shekhar's answer was that it begins with

everyone. Treating security as someone else's job, he suggested, is among the first and most consequential mistakes an organization can make. Partly for that reason, Caltrans has avoided **open-source models** out of concern over malicious prompt injection, running only fine-tuned models within its own cloud environment so that the tuning does not flow back to a public model.

Another participant raised the problem of **legacy data**: agencies will not discard decades of records and applications, so how should that history coexist with new collection? The Caltrans response was to triage: there is rarely a need to translate data from the 1920s or 1950s, and the practical task is to identify the recent five, ten, or twenty years a model actually needs. Shekhar illustrated the point with the observation that Caltrans holds data reaching back as far as 1912, and that the depth of a record materially changes the confidence a model can place in its output. He acknowledged that the agency still lacks a strong enterprise tool for data lineage and age and named acquiring one as a next step.

On **copyright**, he described signing a memorandum of understanding with AASHTO governing the use of AASHTO manuals; crucially, the tool ingests those manuals but does not train on them and runs in a disconnected environment that keeps an arm's-length separation between the agency's content and the underlying model.

Asked whether new **staff positions** were being created, he described two broad categories of emerging role: strategy roles and delivery roles. Pressed on whether agencies need new technical specialists such as data scientists or ethicists, he reframed the question around professional accountability: a licensed engineer signs off on the work product regardless of whether AI was involved, and the test is whether that engineer is comfortable putting a name to the output. On the anxiety that AI will displace staff, he pointed to recent industry analysis suggesting the effects are a mixed bag, falling unevenly and not always where predicted, and argued that the agency's task is to write a counter-narrative in which AI changes work rather than simply eliminating it.

Caltrans also offered an early, concrete workforce example: the agency had begun using AI-driven "agentic" communication on a professional networking platform to engage candidates for hard-to-fill engineering positions, answering questions about base versus total compensation, benefits, and accommodations, an application Shekhar believed put Caltrans among the first public agencies to deploy the technique. Closing the session, Riley added a caution that would recur: even with retrieval-augmented systems that ground a model in agency data, model drift is real, and it demands a defined validation cycle, a discipline, he noted, not normally considered in the TSMO world.

Caltrans AI Projects

Introduction to the Conversation

Following the vision, the Caltrans team turned to the specific tools it has put into production. Rajpreet Khangura of Caltrans presented the agency's Road Safety Insights work, and Jaswinder Gill of Caltrans presented the Traffic Mobility Insights platform. Together their accounts traced a path from a gubernatorial directive to two live, award-winning systems in less than two years.

From Executive Order to Production

The work began with Governor Newsom's Executive Order N-12-23 in September 2023, which directed state agencies to explore generative AI. Caltrans' research division led the effort and ultimately ran two generative-AI initiatives in parallel. The procurement model is itself instructive: rather than write a fully specified contract, Caltrans used a performance-based approach that described the bridge it needed to cross from one point to another and asked vendors how they would build it. It let proof-of-concept contracts for one dollar each in April 2024, completed those proofs of concept by the end of the year, and selected the vendors and awarded the work in early 2025. Two production tools resulted, developed on parallel tracks with separate contractors: one focused on safety, the other on mobility.

The Traffic Mobility Insights platform went live on August 25, 2025, and the Road Safety Insights platform went live on November 1, 2025. The effort has since drawn outside recognition: the California Transportation Foundation named the Road Safety Insights work a Safety Project of the Year, and both tools earned internal recognition.

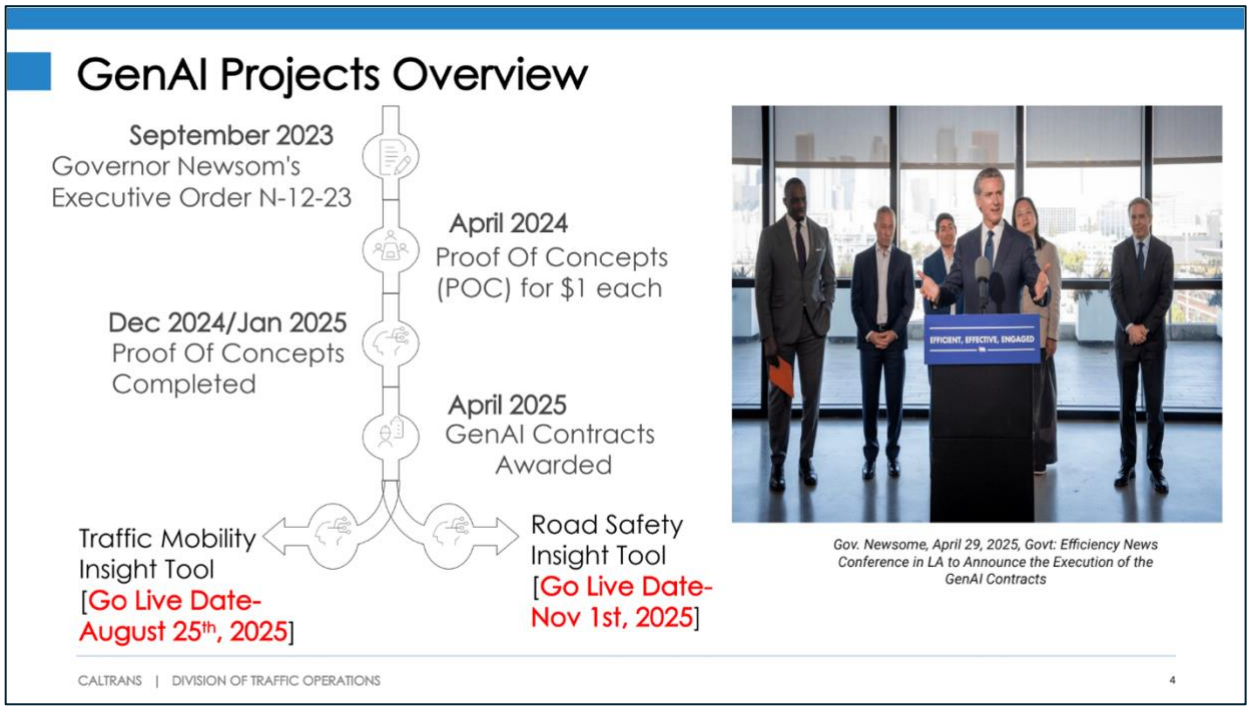

Figure 2. The Caltrans generative-AI timeline, from Executive Order N-12-23 in September 2023 through one-dollar proofs of concept to two production tools live in 2025. Source: Caltrans presentation.

The team was careful to distinguish these purpose-built tools from general productivity assistants. District engineers, Gill noted, kept confusing them with the office productivity assistant the agency also uses, and he built a slide specifically to draw the line. Where a general assistant offers broad productivity support integrated across office applications, the Caltrans tools are tailored to specific program workflows, run against the agency's own traffic-volume and crash datasets and its policy and standards documents, including the California Manual on Uniform Traffic Control Devices, and generate outputs in the formats the programs actually use.

Road Safety Insights

The Road Safety Insights platform began as a generative-AI tool focused on the safety of vulnerable road users (pedestrians and cyclists) and has since expanded to address all crash types. Through a conversational assistant layered over a map-and-data interface, the platform lets a user generate a vulnerable-road-user safety analysis report, with countermeasure recommendations, in minutes rather than weeks. A user can then prioritize corridors, ramps, and intersections on the State Highway System for infrastructure enhancements, surface the systemic factors behind crash patterns. One use case prioritizes transit-stop enhancements with explicit attention to underserved communities and to factors such as lighting, weather, and transit access. The platform draws on roughly thirty datasets: crash and incident records, pedestrian and bicycle volumes, demographics, land use, and a decade of asset-management and project-pipeline data. It has grown from an initial scope covering only a fraction of the data, focused on intersections and vulnerable users, to one that now incorporates all crashes statewide, including segments, ramps, run-off-road events, and wrong-way driving. As the tool broadened, the team added a separate systemic score incorporating vehicle crashes and a filter that lets users switch between an all-users view and a vulnerable-users-only view.

Khangura emphasized that the value of such a tool is to multiply the expertise of the staff who use it rather than to replace them, to give a safety engineer instant access to expert knowledge and analysis. He was equally direct about the discipline the tools require. Realizing the benefit operationally, he noted, depends on prompting skill, and Caltrans has begun developing internal prompting instruction; each district has also been asked to identify at least five candidate projects and to re-examine its existing workflows for places where the tools can help. Caltrans restricts access to crash data by district, so staff see the records for their own area rather than the entire state, and the platform underwent three weeks of formal user-acceptance testing with subject-matter experts before launch.

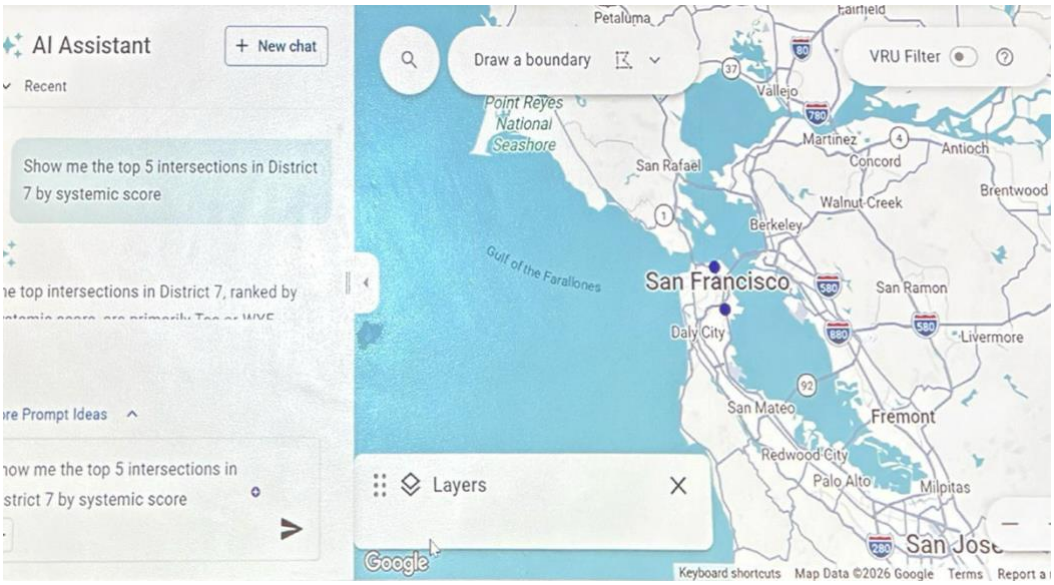

Figure 3. The Road Safety Insights platform as demonstrated live: a conversational AI assistant (left) answering a question about high-priority intersections, alongside the map, boundary-drawing, and data-layer controls operators use. Source: Caltrans live demonstration.

Traffic Mobility Insights

Jaswinder Gill described the Traffic Mobility Insights platform as Caltrans' first generative-AI solution in production for mobility management. Built with a private-sector partner and running on a secure cloud foundation, the platform integrates seven statewide systems into a single geospatial environment and draws on real-time volume data from tens of thousands of roadway sensors, processing very large data volumes at low latency. Data sources include the state's performance-measurement system, its crash-reporting system, and California Highway Patrol incident data. The recommendation engine draws on a corpus of more than 1,200 reference documents from Caltrans, the Transportation Research Board, FHWA, AASHTO, and the federal Manual on Uniform Traffic Control Devices. The vendor cleans and normalizes data through a tiered "medallion" architecture before the tool uses it, and an early lesson was the need to teach the system that a single road may carry several names. For example, a coastal highway might appear as a numbered route, a state highway, or a local street name, and the system must recognize all of these as the same roadway so that a plain-language query will produce accurate results.

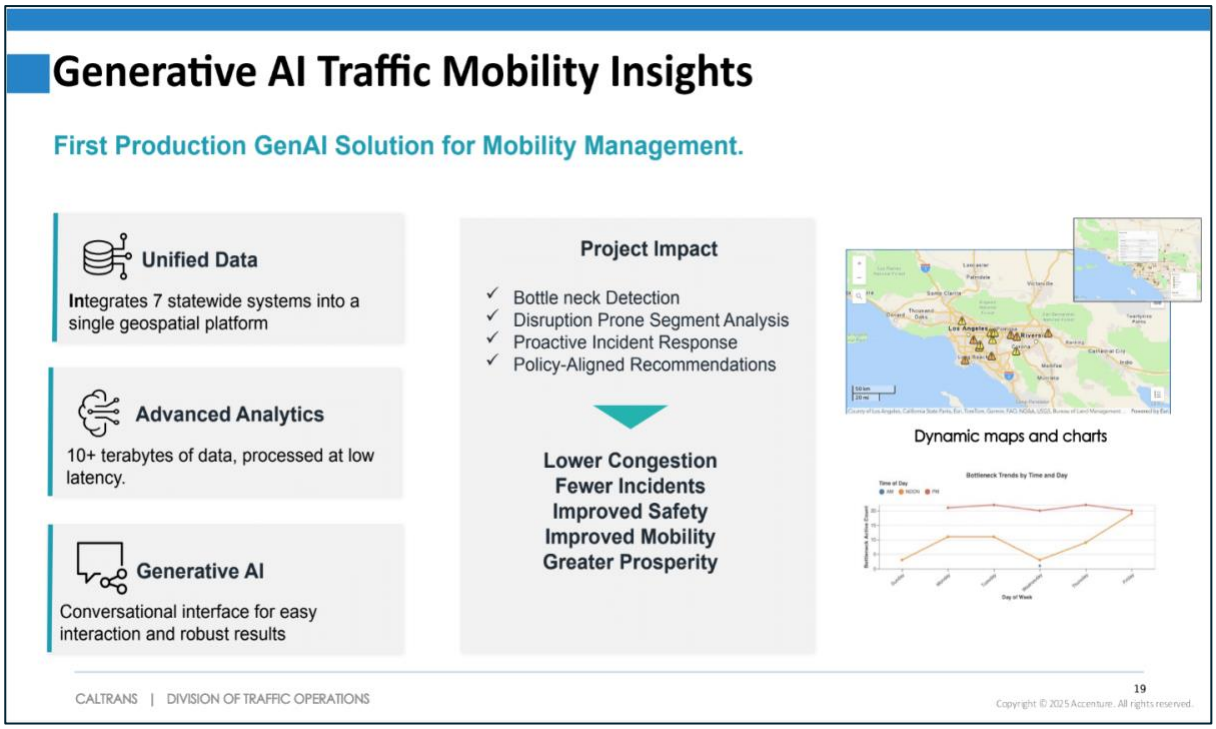

Figure 4. Traffic Mobility Insights at a glance: seven statewide systems unified into one geospatial platform, more than ten terabytes of data, a conversational interface, and the outcomes the program targets. Source: Caltrans presentation.

Through a conversational interface, operators can identify recurring bottlenecks on a given segment, detect incidents or unusual patterns at a location, and access potential countermeasures drawn from that document repository. In a live demonstration, the platform pulled the most recent crashes and incidents from the Highway Patrol feed and let the user ask follow-up questions about active work zones or weather. An unanticipated benefit emerged in use: because the tool reads the sensor network continuously, it flags broken or non-reporting sensors by location, producing a maintenance to-do list as a byproduct. One efficiency gain captured the value directly: the quarterly congestion report districts assemble for FHWA was built by hand across roughly fifteen spreadsheet tabs in a task that consumed somewhere between ten and forty hours. It now can be generated in under three minutes and arrives very nearly ready to use. Every response the platform produces carries a standing disclaimer: "The content is AI-Generated, please be sure to cross check," which is a small but concrete artifact of human oversight built into the tool itself.

Gill was asked how long it took before the team trusted the platform enough to rely on it. The answer was roughly six months, moving from a proof of concept built on static data, through testing, to a light production

deployment. Shekhar added a standard the team holds to: **for a tool to be worth deploying, its accuracy must exceed what a person would produce unaided**. The team was also candid about cost. Cloud hosting fees accrue whether or not anyone uses the tool, so the efficiency the tool produces has to be real enough to justify the bill, and the program is increasingly tying its use cases to existing project identifiers to demonstrate that value and sustain funding. Asked about collaboration beyond Caltrans and the roadways the agency manages, the team described a vision of extending access to other California public-sector agencies, metropolitan planning organizations and transit agencies among them, while acknowledging open questions about additional data sources, user security, and cost-sharing that the agency is still working through.

Gill closed with the lessons the two projects had taught, and they reinforced the program's data-first conviction. The team urged agencies to:

- define their data needs precisely and identify the owners of each data source
- establish secure data pipelines
- validate data and scrub out personally identifiable information before loading it
- map existing program workflows and find the automation opportunities within them
- develop in coordination with IT and within existing IT standards
- plan from the start for maintenance and operations sustainability
- address cloud vulnerabilities with mandatory role-based access controls

Caltrans considers itself early on the adoption curve. The agency's implementation path runs from early adopters among headquarters and district subject-matter experts, through daily-workflow adoption, to incremental enhancement and expansion into new use cases. Its research division continues related work, including developing evaluation criteria for machine-learning and third party-generated traffic data and exploring AI agents to automate tasks in traffic management centers.

AI in Pooled Fund Studies

Introduction to the Conversation

The morning's final session widened the lens from a single agency to **multi-state models**. The Transportation Pooled Fund (TPF) Program is a collaborative program administered by FHWA in coordination with State DOTs. The program allows agencies to collaborate with peers to combine their monies, subject matter expertise, and resources to perform high-priority research to address a diverse range of shared transportation issues. Marlon Spinks of Michigan DOT presented the work of the long-running ENTERPRISE pooled fund study, and Giri Venkiteela of New Jersey DOT presented a newer pooled fund study effort organized specifically around AI. Ramfos framed the pairing around a simple question: What does a multi-agency model make possible that no single state can do on its own?

The ENTERPRISE Pooled Fund

ENTERPRISE TPF-5 (490) is a long-standing multi-state pooled fund study, underway since 1991 now in phase III, that evaluates new technologies for the benefit of its member agencies. Michigan DOT is the lead agency, and the study chair works for Michigan DOT. Its most recent product, which Spinks presented, is a final report on the role of artificial intelligence in intelligent transportation systems (ITS). The report is a practical reference rather than a research treatise and Spinks described it deliberately as an introductory starting point rather than a prescriptive standard. It lays out considerations for agency AI policies, among them

interpretability and transparency, accountability and human review, privacy and confidentiality, fairness, and ongoing assessment. It lays out considerations for implementing AI in operations, including data quality and the risk of over-saturating a model with more data than it needs, workforce and training needs, institutional policy and risk management, computational resources, and the trade-offs between a pilot and a full-scale deployment. It also catalogs example use cases by domain: signal operations, highway operations, road weather management, and asset management, alongside a compilation of national AI resources, state DOT AI policies, and enacted state AI legislation. The report is available at https://enterprise.prog.org.

Spinks grounded the discussion in a taxonomy drawn from US DOT's ITS Joint Program Office, which describes how AI can replace or augment the work of field sensors, connected and automated vehicles, traffic management center operators by identifying objects and images, recognizing speech and audio, processing large volumes of data to find patterns, learning from experience, and providing situational awareness. He pointed to a concrete example from a neighboring state that uses in-vehicle camera imagery to detect roadway weather conditions and noted a companion 2026 research project to adapt vehicle "glass" cameras to detect roadside assets. Drawing on Michigan's own experience, he offered one caution that would recur in the contracting discussion: agencies should not assume that training data effective in one state will produce the same results in another.

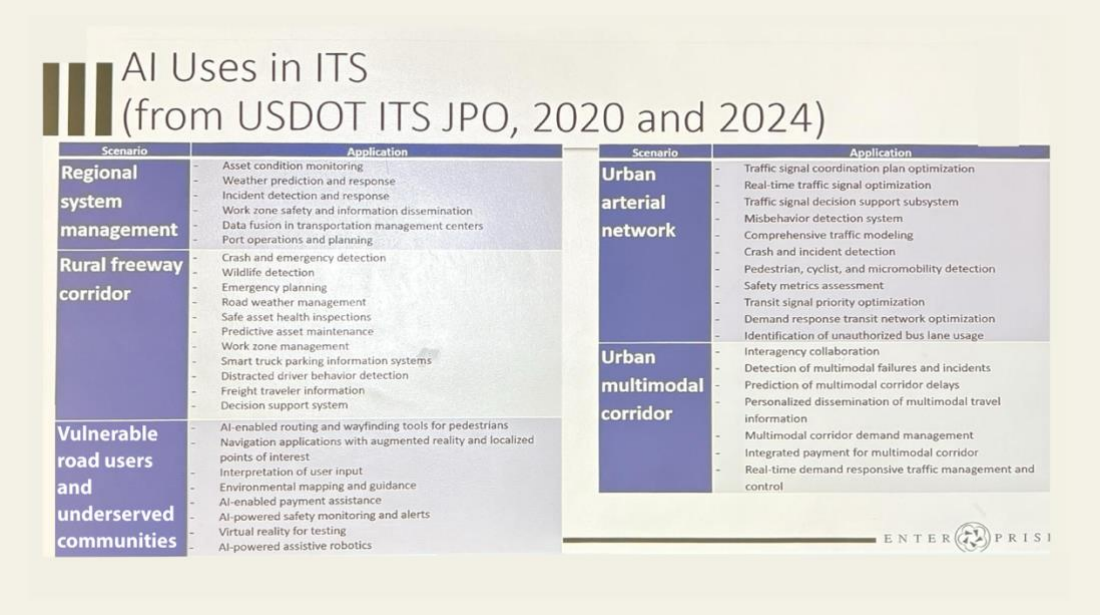

Figure 5. Where AI shows up in intelligent transportation systems: the US DOT ITS Joint Program Office taxonomy presented through the ENTERPRISE pooled fund, mapping applications by setting from regional system management to urban multimodal corridors. Source: Michigan DOT presentation.

New Jersey DOT and the State DOT Council for Strategic AI Adoption

Giri Venkiteela presented a newly established and more focused research vehicle: a transportation pooled fund study, designated TPF-5 (573) and led by New Jersey DOT, established to create a State DOT Council for Strategic AI Adoption. The idea originated about two years earlier among state DOTs meeting in San Diego, with the premise that fifty states should not each have to solve the same problem fifty separate times. The study recently received clearance from FHWA and attracted fifteen participating agencies at the time of the Peer Exchange: California, Connecticut, Kentucky, Massachusetts, the Maryland State Highway Administration, Minnesota, North Carolina, New Jersey, Oklahoma, Texas, Utah, Vermont, Virginia, Washington, and Wisconsin. The contribution structure is modest and deliberately scaled: roughly $20,000 a

year for two participating members, or $10,000 for one, across a five-year project, with total commitments of $1,225,000 received to date.

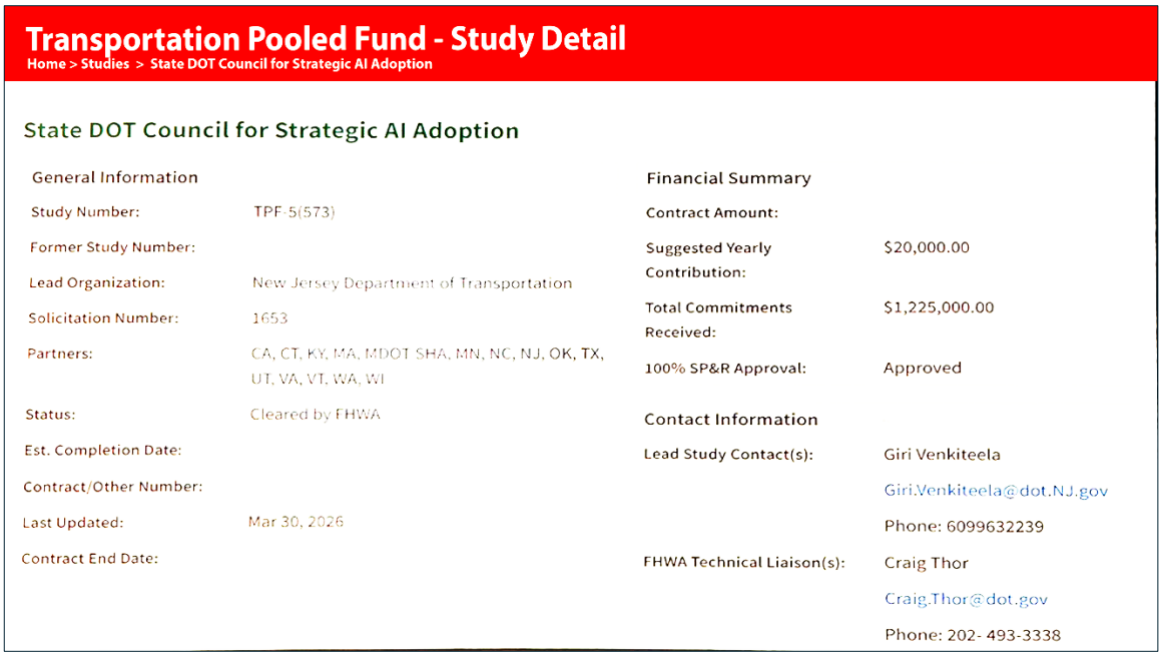

Figure 6. The State DOT Council for Strategic AI Adoption as posted in the federal Transportation Pooled Fund system — study TPF-5(573), New Jersey DOT lead, fifteen partner agencies, and $1.225 million in commitments. Source: New Jersey DOT presentation.

The State DOT's AI Council's purpose is to turn shared cost into shared capability. The scope includes:

- organizing biannual peer exchanges for member states to compare AI implementation experiences
- developing actionable deployment strategies
- running multi-state pilot projects on common use cases to validate technologies and demonstrate best practices
- standardizing and improving data quality across participating DOTs to produce AI-ready datasets
- creating shared resources, including research repositories, an implementation toolkit for states newer to the work, and workforce-training materials.

Venkiteela described the early work as a period of knowledge exchange, establishing who holds what data and who oversees it. Then the work will move to shared pilots, shared datasets, and shared training, with findings fed upward to the federal level. He laid out a concrete near-term schedule: a request for proposals to facilitate the peer exchanges in May 2026, steering-committee review in June, contract onboarding in July, a virtual kickoff in August, and a first in-person peer exchange in Dallas, Texas, on a date to be determined.

The discussion provided clear insight into the logic of the pooled fund study model. The central advantage participants named was the distribution of financial risk; a pooled approach lets agencies attempt work that no one of them could justify funding alone, and it is especially valuable to smaller states. The representative from the Vermont Agency of Transportation, the smallest member, said plainly that a shared investment of that size was significant for an agency that could not run forward on its own. The council also offered a direct answer to a thread that had run through the morning: it is the natural place for an agency such as Caltrans to share the large inventory of use cases it has developed, so that one agency's work can spark another's.

Partner Initiatives

Introduction to the Conversation

The afternoon gave the national associations and federal partners a structured opportunity to describe what they offer DOTs and, as importantly, what they need from them. Five organizations presented in turn: AASHTO, represented by Ahnaf Morshed; ITE, represented by Ted Bailey of STV, who also serves on the NOCoE Technical Advisory Committee; ITS America, represented by Charnita Allen from the District Department of Transportation (DDOT), and Kevin Viita from ITS America; and US DOT, represented by Joseph Gregory.

AASHTO

Ahnaf Morshed presented the results of AASHTO's national survey of state AI activity, which had been previewed at the 2025 peer exchange and now formed the most detailed picture of the field available to the room. AASHTO launched the survey in March 2025, which was fielded over roughly three and a half months, and drew about seventy-six responses spanning every region, consisting of nineteen questions. Respondents were a mix of leadership and hands-on staff, with technology executives and director-level staff together making up well over half of the participants. Rather than a binary measure of whether an agency had "adopted" AI, the survey built an AI coverage map from three indicators: familiarity, preparedness or readiness, and deployment.

The findings revealed a field that is more familiar with AI than it is ready for it. Most regions described themselves as familiar or somewhat familiar, but few as genuinely prepared, a sign that the technology is running ahead of organizational readiness. Asked to rate the importance of various application areas on a five-point scale, respondents placed traffic management and operations highest, with safety and incident management and data analysis and decision-making close behind. The leading challenges were data quality and availability, cited by about three-quarters of respondents, and cybersecurity, cited by roughly seventy percent, followed by trust and reliability, workforce, and regulatory and legal barriers. About twenty states reported having or developing an AI strategic plan. About a third had already uploaded their specifications and standards into AI tools, which Morshed flagged as a concern when those tools sit in closed commercial ecosystems. AASHTO has published a white paper summarizing the survey, maintains a live dashboard of the results, and was preparing a second version of the survey that would add questions on responsible AI and workforce impact, along with a maturity framework developed with partners.

ITE

Ted Bailey introduced the practitioner-led work of ITE's AI in Transportation community, framing it around the question of workforce capability rather than tool enthusiasm. He emphasized how ITE is a connector to the local-agency world: only about eleven percent of ITE's members are state DOT staff, while forty to forty-five percent are local-agency public employees. An ITE practitioner survey had found that only about eight percent of respondents considered practitioners "ready," from a small sample of roughly a hundred out of some nineteen thousand members, a base the group hopes to grow into the thousands. Even that early sample mapped twenty to thirty AI tools already in use, and one respondent estimated spending around four hundred dollars a month in personal subscriptions to keep a working toolbox. ITE's products include an annual survey with plans for a live dashboard, a self-assessment kit, an effort to document how a task was performed before and after AI was introduced, and a student innovation challenge. Ted Bailey's emphasis on a structured readiness assessment that agencies and individuals can use to locate themselves, anchored the Day 2 workforce session and is described in full later in this report.

ITS America

Charnita Allen offered the perspective of an urban agency, the District Department of Transportation in Washington, DC, operating inside a centralized municipal IT environment. The District requires staff to complete three courses before they may use AI tools: an introduction to AI, an introduction to generative AI, and a course on the risks and policies of AI. Staff use of the standard productivity assistant, run short "tech talks" to spread practical knowledge, and must document and obtain approval for any tool that contains an AI component. The account was a useful counterpoint to the state DOT experience: a smaller, denser jurisdiction with centralized technology services faces a different set of constraints and a different path to adoption.

Kevin Viita presented a bundle of ready-to-use resources and offered them explicitly for inclusion in this report. ITS America's AI committee, he noted, is chaired by Nathan Wade of Flow Labs and Karen Lightman of Carnegie Mellon University. Working with an outside consulting partner, the committee has produced a practical guide to next steps for AI implementation; a self-assessment worksheet, framed as a "path to AI maturity," designed to be worked through by a cross-functional team in a single sitting and to produce a twelve-month action roadmap; a set of AI governance group guides addressing the executive, delivery, and operational functions within an organization; and a course on artificial intelligence in transportation offered through the ITS Academy. The organization was also developing a guide to integrating privacy into AI for intelligent transportation systems, aimed at policymakers and agency executives, which enumerates the privacy risks that accompany large-scale data collection. This includes data misuse and repurposing, a lack of transparency and consent, algorithmic bias, increased surveillance, and identity theft. Beyond AI, ITS America maintains a broader technology use-case library and work on transportation data interoperability, and its annual meeting the following month in Detroit would feature workshops on AI, data interoperability, and cybersecurity. ITS America also has a workforce-development community of practice that helps transportation professionals prepare for emerging technologies, including AI, through workforce-readiness policies and investments. Links to these resources appear in the Resources section.

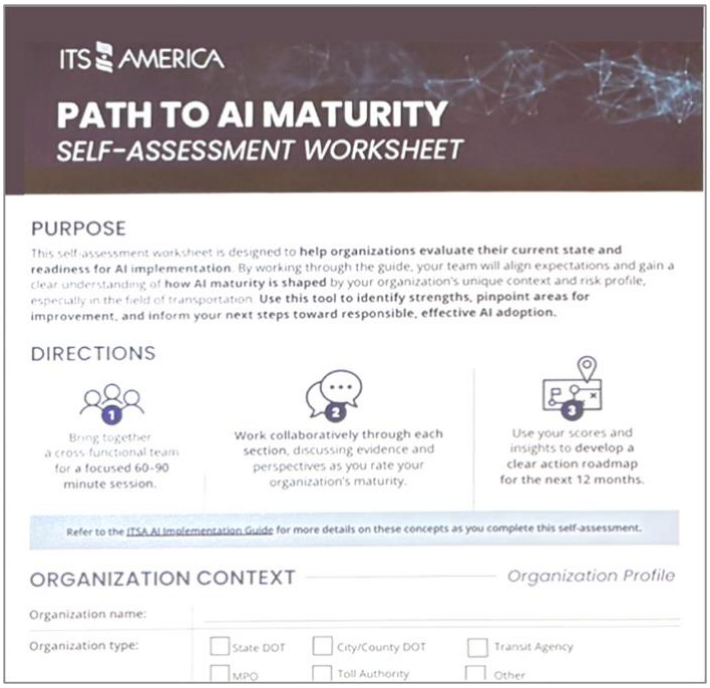

Figure 7. ITS America's "Path to AI Maturity" self-assessment worksheet — a facilitated, cross-functional team exercise that produces a twelve-month action roadmap, shared with participants for agency use. Source: ITS America materials shown at the exchange.

US DOT

Joseph Gregory provided the federal perspective. FHWA has announced the establishment of an AI Center of Excellence: an AI deployment service intended to act as a national clearinghouse for playbooks, traffic models, datasets, procurement templates, data standards, workforce and training resources, and governance frameworks, with a public website and a subscription for updates. The center had launched a National AI Transportation Coalition, structured as a transportation pooled fund with an open solicitation, organized around a five-year strategic roadmap. Gregory described two initiatives:

- the **Intersection Safety Systems** initiative applies AI to existing and low-cost infrastructure to predict conflicts and prevent crashes. Structured as a prize competition that ran from 2023 into early 2025, it had produced multi-sensor training and validation data collected at the Turner-Fairbank Highway Research Center, now downloadable through the federal ITS data hub, with an end-to-end prototype planned for later in the year.
- **AI and traffic management centers** is establishing a data benchmark for training and evaluating AI prediction models and FHWA is carrying it as a task within the traffic management center pooled fund study.

A further effort is developing a framework for evaluating the trustworthiness of AI models, for which the federal team was seeking state DOT case studies.

The cross-cutting discussion among participants returned repeatedly to a single request: a clearinghouse or shared use-case library.

DOT Roundtable on the Uses of AI

Introduction to the Conversation

The first day closed with a roundtable in which each agency, in turn, answered three questions:

1. What AI it is using or piloting
2. What its biggest barrier is
3. What it intends to do next.

Rather than recount each agency in sequence, this section synthesizes what emerged. The single most important result is stated first, because it represents the clearest change since 2025.

Funding Has Become the Leading Constraint

Where cost and over-reliance on technology had led the list of concerns in 2025, the agencies in Sacramento named funding as their foremost barrier and they meant something more specific than a general shortage of money. The difficulty is that AI work sits outside the traditional program. Implementation costs and recurring subscription costs both have to be considered and the shift from owned assets to subscription services strains budget structures built around capital expenditure rather than continuing operations. Formula funding does not fully cover operations, maintenance, or TSMO, and even where a pilot can be funded as a capital expense, the ongoing costs that follow are hard to lock down. Several participants connected the problem to the pending surface-transportation reauthorization, observing that while “digital infrastructure” is a welcome category, agencies want funding directed specifically at AI; some noted pointedly that automated vehicles were attracting attention in the reauthorization conversation and asked what automated vehicles had yet delivered for congestion or emergency response.

The economics of the tools themselves reinforced the point. Agencies returned often to the true cost of the tokens that meter usage of commercial models, and to their desire for consistent, reproducible outputs. Texas DOT noted that costs accumulate quickly. Rose Vath of the Vermont Agency of Transportation offered the room’s most contrarian answer: rather than purchase tokens, Vermont runs its own custom-built large language model on hardware two generations old. The software is billed to the user community as storage rather than as metered tokens and as a result is considerably cheaper than a commercial assistant, with principle of getting to ninety percent capability at low cost and investing in the last ten. Another participant described setting up per-project environments specifically so that token costs could be charged back to the teams using them, keeping the tokens drawn against a federal project from being consumed by routine

administrative work. The other side of the funding problem, several agencies observed, is the burden of demonstrating value. Agencies must produce case studies that show cost savings and must explain AI's benefit to the public, a harder sell than the visible, tangible value of concrete and asphalt. One data leader recalled producing, during the pandemic, the figure that some eighty-five percent of the state's residents had stayed home, only to find that the public did not particularly care, a lesson in how hard the storytelling can be.

The Breadth of Activity in the Room

The applications agencies described ranged across the TSMO portfolio. Rick Zygowicz of Missouri DOT described using AI and cameras to **improve safety at the intersection level**. Greg Dyer of Tennessee DOT described an AI decision engine in the agency's Middle Tennessee traffic management center, supporting lane control, variable speed limits, and video detection, and expressed particular interest in optimizing the many signals that are not connected. Seth Daniels of Nevada DOT described work on the **prepositioning of police and incident response** and noted the complication that the agency owns but does not operate or maintain its signals, and that a large share of land in the state is federally controlled. Aiden Burnside of North Carolina DOT named funding, time, and overlapping projects as the agency's barriers and described working with an AI camera vendor to develop **enforcement for high-occupancy-vehicle lanes, tolling, and distracted driving**. He also noted that the state had recently passed legislation permitting AI in tolling, against a national backdrop of privacy concern and active litigation over the practice. Stin Weber of Maricopa County, in Arizona, described deploying a public-facing chatbot to answer citizen questions and the potential of using AI in the county's permit center to ensure initial submittal requirements are met, reducing review tasks that consume staff time. Texas DOT pressed the theme that data must be managed like an asset, cared for, maintained, and protected, because machine learning and AI depend on verifiable, reliable data. Sanhita Lahini of Virginia DOT named funding and the absence of a cohesive AI policy as the agency's constraints, described limiting access to AI tools and requiring courses before use, and raised the practical obstacle that its field controllers cannot communicate directly with the cloud.

Several common threads ran beneath the individual reports. The problem of unconnected signals recurred across Iowa, Tennessee, and Nevada. The tension between operations and IT surfaced repeatedly. David Man of the Metropolitan Transportation Commission, the Bay Area's Metropolitan Planning Organization, reminded the room that much of the local system has no sensors, communications, or computing at all, and that lifting up small and medium agencies is part of the work. The group also debated where AI belongs in the field's maturity frameworks, whether it warrants a capability maturity model of its own or, as several argued by analogy to safety, should instead be a dimension that runs through the AASHTO TSMO capability maturity model, woven into its culture and collaboration elements rather than standing apart. A connected-vehicle data thread surfaced when a participant challenged vehicle manufacturers to share anonymized occupancy data of how many people are in a car as a base layer that would make tolling, high-occupancy-vehicle enforcement, and ramp metering far easier, while acknowledging the constraints of state privacy law. Agencies also traded notes on a fast-moving software market, where vendors were raising prices by thirty to forty percent as they funded their own AI development and pushing more aggressive contract terms, and on the public-records implications of AI, which cuts both ways: the public can now mine the data agencies release more powerfully than ever, even as AI can help an agency answer records requests faster.

Importantly, nearly every agency, whatever its level of maturity, expressed some version of the relief that comes from discovering that peers face the same unknowns.

Day 2 — Wednesday, May 20

The second day opened with a brief recap connecting Day 1's picture of activity to the day's harder questions. If the first day had established what agencies are doing, the second asked how they scale it and sustain it, beginning with two agencies that have moved from isolated pilots toward program-scale adoption.

Expanding the Use of AI in Your DOT

Introduction to the Conversation

Eryn Stone of Iowa DOT and Ben McCulloch of Texas DOT described practical requirements to expand AI within a state DOT, including examples of the pilot projects that failed. Their accounts were among the most concrete of the exchange, and they shared the common compelling observation that the work depends less on the sophistication of the technology than on data, relationships, and the organizational conditions that let good ideas spread.

Iowa DOT

Eryn Stone began with the policy environment, which is the binding constraint in Iowa. The State of Iowa adopted an AI policy in March 2025 that governs all use. The policy permits content generation, research, decision-making support, proofreading, and the preparation of presentations and training, but it prohibits the use of sensitive data, the use of personal AI accounts, the manipulation of information, and any infringement of intellectual property rights.

Figure 8. "Where we started": the March 2025 State of Iowa AI policy — its purpose, allowable uses, and prohibitions — which frames everything Iowa DOT has been able to attempt since. Source: Iowa DOT presentation.

Within that environment, Iowa DOT has pursued a structured AI strategy in five phases — an AI landscape review and orientation; the identification of use cases mapped against the agency's matrix of mission functions and work types; a capability maturity assessment of its data, workforce, governance, and technology; facilitated workshops to refine and finalize the strategy; and the development of a roadmap.

A telling structural fact shapes much of Iowa's traffic signal work: the state DOT owns and maintains zero traffic signals, all of which belong to local agencies whose maintenance practices and technical knowledge vary widely. That is why the agency's signal-operations tool cannot obtain the timing data it would need as cities are unlikely to share it, and most signals are not on fiber networks. Iowa has nonetheless demonstrated genetic-algorithm signal-offset tuning with a human in the loop on a test corridor.

Stone then described several pilots that did not work. A crash-prediction tool built with a vendor used a wide range of datasets: employer records, probe-vehicle data from multiple providers, location data, and weather to predict crashes and their severity by county. But the design offered no feedback loop or verification of its predictions. Another work-zone detection pilot with a vendor, drawing on the work zone data exchange (WZDx) feed to bound the area of interest and testing haul routes before and after, ultimately didn't deliver substantial outcomes for Iowa. A separate pilot used AI to scrub publicly available web data, including permits and 811 utility-locate requests, and to evaluate satellite imagery to identify utilities in the right-of-way; used on two projects, it produced interesting results but ran into gaps in data availability that prevented full validation. On the more promising side, two Johns Hopkins University students, given open agency data, built a prototype AI phone service to route callers to the correct maintenance garage and report road issues. This tool is not in production but was a demonstration of what is possible.

What Iowa DOT uses in production today is concrete and operational: electronic ticketing through a construction-logistics platform, AI-generated daily work summaries, connected contractor equipment, active-jobsite data drawn from a navigation app, and work-zone monitoring. The agency has also put computer vision to work on its camera network, scanning the full video stream to identify weather conditions and flag obstructed views. The agency's forward strategy reflects a deliberate choice that distinguishes it from peers with standalone plans: rather than create a separate AI strategic plan, Iowa intends to fold AI into its existing data management strategic plan and to designate standards and budgets for each functional category. Iowa also pointed to a multi-state effort, the National Operations Digital Environment, as a platform for sharing work-zone, incident, detour, and condition data across state lines along major corridors.

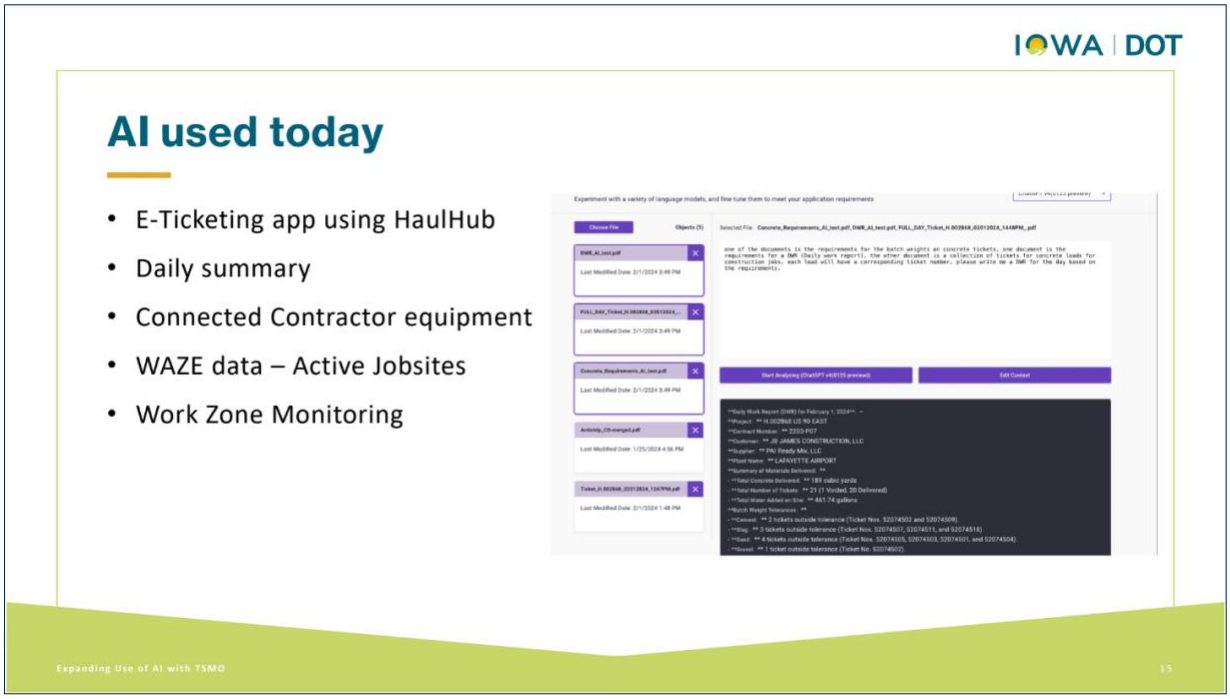

Figure 9. AI in production at Iowa DOT today: electronic ticketing, AI-generated daily work reports, connected contractor equipment, navigation-app jobsite data, and work-zone monitoring — including a demonstration of a GPT-4-based tool assembling a daily work report from raw delivery tickets. Source: Iowa DOT presentation.

Texas DOT

Ben McCulloch provided a full account of Texas DOT's journey, tracing the entire arc from incompatible spreadsheets to an agentic AI system, making the case that the technology is the easy part. McCulloch, a data scientist at Texas DOT, located the agency's progress not in any model but in relationships and goodwill, built patiently over roughly eight years. Trust, he argued, was what eventually allowed his team to move quickly.

The work began well before the release of AI tools, with what he called the problem of building bridges between islands of data. Texas DOT, like most large agencies, carried several incompatible ways of describing the same point on the road: mileposts, reference markers, latitude and longitude, control sections with no reliable crosswalk between them, so that joining two datasets was a manual ordeal. His own start at the agency was modest: Texas DOT hired him around the end of the last decade to maintain a single dashboard, at a time when only a handful of people-built dashboards and perhaps fifty of the agency's twelve thousand staff had access to them. By establishing standards and training citizen developers, the agency grew that to some fifty trained builders and, in McCulloch's phrase, opened the floodgates. The breakthrough that converted skeptics was unglamorous: McCulloch's team re-engineered a weekly financial-tracking process spanning roughly a hundred and twenty categories of money, one that had taken one person eight hours every Monday to run automatically six times a day. The IT organization had quoted a year and a half to do it; his team did it in about eight months, and the finance office went from the effort's biggest doubter to its biggest champion. Another example showed the safety stakes of the same problem: by finally linking maintenance, crash, and project data, the team discovered that a curve had been inadvertently altered during a resurfacing project and was generating run-off-road crashes, a pattern that had gone undetected for months only because the data had been too hard to pull together.

That foundation is what made the agency's AI work possible. Texas DOT built an **AI strategy** through fourteen focus groups and ninety-one participants, drawing in all twenty-five districts and its divisions, generating some five hundred sticky notes and more than two hundred candidate use cases, and it now prioritizes those use cases through a scoring rubric. The agency published its first **AI strategic plan** and issued an update in early 2026. Underneath sits an enterprise data platform built on a commercial data-cloud, integrating roughly thirty systems with spatial data added in 2025, and structured so that districts and divisions can write their own queries and promote work through development, test, and production stages rather than waiting on IT. A sandbox built on a major cloud provider's machine-learning and model services gives staff a place to experiment, and the whole environment is built on a retrieval-augmented design with no open web access. The models see only the agency's crash data, manuals, and data dictionaries.

The marquee demonstration was an agentic generative-AI tool the team had nicknamed DAVE, for its function as a data analytics and visualization expert. Rather than a single model, it is an orchestrator coordinating a set of specialist sub-agents, each an expert in one dataset: crashes, pavement, projects, and weather. As a compiler agent, the tool assembles and returns to the user a map in the agency's geographic information system. Each sub-agent's instructions live in a separate file under version control, and because the team had pre-aligned the underlying data to control sections and distance-from-origin, the system does not have to reconcile competing location references on the fly. The tool keeps session memory; it does not yet reach real-time data, which McCulloch said the team deliberately held back for a future agent. McCulloch quantified the value proposition: a query of the kind the tool answers in moments had typically required three or four people working a full day each, on the order of twenty-four to thirty-six staff-hours, and, against an estimated five thousand uses and a loaded labor rate of roughly fifty-six dollars an hour, the team projected savings somewhere between $6.7 million and $10.1 million.

McCulloch's account doubled as an argument about data discipline, and it produced the line that lingered longest in the room: Free text, is the bane of your existence. The agency's bridge system alone held some two hundred and fifty different ways of writing "Interstate 35," and the plan was to use AI itself to build a crosswalk from each variant back to a standard form, recycling work another group had already begun, after an earlier standardization push had died amid managerial infighting. He was frank about the fragility of the capability as well: the team responsible for all of this had shrunk from two people to one. He closed on a note

of urgency, describing the present moment as a once-in-a-generation opportunity to get the agency's data house in order.

Questions and discussion. The discussion that followed was extended, and this report captures it in some detail because the questions drew out how Texas DOT had actually managed the change.

Asked how he had been allowed to let his team build all of this, McCulloch returned to his central theme: there had been no single authorization, only relationships and goodwill accumulated over eight years that made people willing to say yes.

A participant from Missouri DOT, describing his own agency's chatbots as tools that users tend to treat like a text-message exchange, asked whether Texas DOT was letting its tool ask questions back to the user, and how that was going; McCulloch confirmed that the team was indeed enabling the system to ask clarifying questions, but within guardrails — the tool checks whether a question is engineering-related and confines the geographic scope of a query, for instance to a fifty-mile radius rather than letting it range across the whole state. Asked whether the system caches data, he said it caches neither prompts nor answers.

Asked whether any of the data could be made available to the outside world, he said it could not, though some data is published through dashboards.

Asked how an agency gets its people to want to use these tools and to invest in learning them, he pointed to the work of bridging the business and IT sides of the house and simply getting projects moving so that staff can see results.

Asked about the size of his team and whether universities or research centers were helping, he gave the candid answer that the team had gone from two people to one.

Rose Vath from the Vermont Agency of Transportation asked how he evaluates the tool's outputs, whether he checks for accuracy or for completeness. McCulloch's answer was that the team checks for accuracy and probably for completeness as well, but only for querying the data, not for engineering decisions, which remain with the engineers.

The morning's broader discussion took stock of where agencies stood. Only a handful reported having an AI strategic plan in hand: Texas, the District of Columbia, California, and Nevada among them, with several of those still in draft and Virginia working through a council. Participants identified the recurring tension between operations and IT explicitly, and one concrete request surfaced for AASHTO or NOCoE: a best-practice study of how ITS and TSMO should be organized within a DOT, addressing contract authority, funding authority, and the idea of the agency acting as the product owner of its own systems. Ramfos noted that NOCoE's peer exchange on the boundary between information technology and operational technology, held the prior year, and its published report begin to answer those questions.

AI in Contracting

Introduction to the Conversation

The contracting session was an all-attendees working discussion, grounded in NOCoE's prior peer exchanges on TSMO and information technology, on AI for TSMO, and on innovative procurement. A bank of prepared questions structured the conversation, but the facilitators went deep on a subset rather than attempting to cover each one. What emerged was less a set of answers than a clear map of where the unresolved problems lie.

Where the Problems Lie

The first problem is that AI increasingly arrives hidden inside other procurements. Vendors now embed AI capabilities in tools agencies already buy: advanced traffic management systems, traffic-signal controllers, video analytics, and the broad run of software-as-a-service products. Agencies struggle to identify when AI is a component of a purchase and whether to treat it differently. Caltrans described handling this through careful contract wording. The harder distinction, several agencies noted, is between specialized AI such as machine learning and generative AI, because policy often addresses only the latter. California's statewide AI procurement policy, set through the state's technology department, is focused on generative AI, which leaves agencies uncertain whether more specialized machine-learning work, a simple visual-perception model, must run through the same assessment gate as a large generative system, and uncertain whether AI embedded in something like a learning-management or training package should be evaluated at all.

A second cluster of problems concerns how to vet what is being bought. Texas DOT described a risk-management working group that must review and score any proposed AI tool as a built-in step in procurement, and an emerging-technology pilot program, run with a university partner, that tests whether a vendor's tool performs before it is scaled. A recurring constraint is the prohibition on sole-source procurement: open solicitations extend timelines but make eventual deployment easier to scale. Agencies also acknowledged, when asked directly, that most have no formal definition of "AI," which complicates the writing of enforceable requirements. And the mismatch between technology and contract vehicles drew specific examples including cellular-connected traffic cameras, for instance, may belong under a wireless procurement vehicle rather than an IT one, and where a tool lands can be an accident of the industry classification codes that procurement systems are organized around. Ramfos pressed two further questions the room could not close. The first concerned cooperative purchasing: agencies have moved faster on ITS technology through national purchasing alliances and multi-state agreements, and whether those vehicles work for AI tools and services remains an open question. The second concerned performance metrics: what does "it works" actually mean in a contract for a probabilistic tool? The most defensible answer, offered by Govind Vadakpat of FHWA, was that it depends on the use case; the room had no general standard, and participants noted the absence. Several agencies also flagged a two-tier accountability gap created by systems integrators, who specify equipment and effectively deliver a second product, the manufacturer's embedded AI, with its own specifications and warranties that the integrator may not fully understand.

A third set of problems concerns governance, data, and accountability. New York operates an AI review board through which any project involving AI must pass; Oregon has added an AI review layer; Michigan DOT runs its AI procurement through the same technology-and-budget review and prioritization that governs all IT projects, with purchases under a set dollar threshold able to proceed and larger ones going to a transportation board that meets only twice a year, a real timing constraint, and one the agency was discussing how to relax. Vermont handles AI procurement through its IT partner but flagged a deeper gap: engineers with no product-management discipline often write the contracts, and the central IT agency that inserts requirements is not present after award to validate that they were met, leaving the DOT, in one participant's words, simply paying invoices. When an agency trains a model on its own data, it typically builds in layered protections: Virginia DOT maintains multiple layers of protection, and routes copyright questions to the state Attorney General's office. Caltrans identifies its data sources and their limitations and runs only fine-

tuned models within its own cloud so that the tuning does not flow back to a public model and maintains an intellectual-property governance group.

On liability, the consensus framing was that the human remains responsible: an agency cannot offload accountability to a software model, and that responsibility extends to contractors acting on the agency's behalf. The legal reality several participants invoked is that an agency must be able to defend what an AI-assisted decision recommended, and the assertion that "AI told me to do it" will not hold up in court. More than one participant noted the tort exposure created when AI surfaces a known hazard that the agency then fails to act on. The practical safeguard the group converged on is the distinction between a decision-support system and an automated decision system: AI should inform a knowledgeable human's judgment, and the review that follows must be genuine rather than a rubber stamp. Caltrans described the discipline concretely, its safety tool deliberately produces a "digest" rather than a "report," covering roughly 55 to 70 percent of the workflow and leaving validation and sign-off to the safety engineers. The system stops short of becoming an automated decision system, and Caltrans recounted having to clarify that an automated, still-analog traffic-signal system is not an "automated decision system" in the sense the policy contemplates. One cost comparison captured the stakes vividly: a corridor study on a state route weighed roughly $500,000 for AI-enabled cameras against several million dollars for physical roundabout construction to address the same travel-time problem.

Workforce and Training

Introduction to the Conversation

The afternoon turned to the human side of AI, on the premise that training and workforce capability are as essential as the technology itself. Vidhu Shekhar of Caltrans returned to describe his agency's workforce-adoption strategy, and Ted Bailey of STV presented a readiness assessment developed through ITE. Both treated organizational change and data literacy as strategic challenges rather than technical afterthoughts, and the discussion that followed opened the workforce questions about unions, job classification, and the talent pipeline that may prove the hardest of all.

Caltrans' Workforce Adoption Strategy

Shekhar was emphatic that the work is a data *and* AI strategy, run on two parallel tracks: data literacy and data governance first, then AI literacy and AI governance, on the principle that one cannot put the cart before the horse. Caltrans ties the strategy to the agency's broader goals of empowering a people-first workforce, delivering measurable efficiency, and governing AI responsibly, and aligns it to the U.S. Department of Labor's AI literacy framework, whose core competencies are understanding AI principles, exploring AI uses, directing AI efficiently, evaluating AI outputs, and using AI responsibly. The alignment runs upward as well as downward: the governor's office had signed a memorandum of understanding with several major AI providers, focused on workforce, and Caltrans was cascading that agency-level agreement down to the department and coordinating with the state labor agency, so that the effort joins from the bottom up and the top down at once.

The strategy answers a question agencies repeatedly named as difficult: how to measure adoption. Caltrans measures it across five lenses: readiness, adoption, capability, value, and trust where trust is understood in two dimensions, the workforce's internal trust in the agency's data and systems and the public's trust in the agency itself. The agency had given pilot licenses for its productivity assistant to roughly seventeen thousand of its twenty-two thousand employees, excluding field staff without email accounts, and it tracks adoption through a distribution curve of how many people are using the tool and how often. Crucially, Shekhar stressed, the program watches that curve at the organizational level and never at the individual level. The goal is to shift the whole curve upward over time so that fewer people sit at zero, not to single anyone out and the

tracking is specific to the productivity platform rather than to token cost, which he described as an all-you-can-eat arrangement rather than a meter.

The implementation rests on four pillars: a tiered set of training pathways, centers of excellence aligned to platforms, communities of practice, and peer-to-peer learning. Training runs as a four-tier pyramid, with a basic awareness tier mandatory for everyone and progressively smaller tiers above it, and the agency uses generative AI within its learning system to generate dynamic quizzes, so capability is in fact tested, not merely assumed. Staff must periodically certify that they remain responsible for their own work product regardless of any AI involvement. Rather than build new communities from scratch, Caltrans is leaning on existing ones, the first community of practice is in construction, and a flagship "communications collective" of communications professionals from headquarters and every district will present a "day in the life" of their work as enhanced, and disrupted, by AI. The agency also found that its vendor's standard training was not good enough and sent it back to be rebuilt around personas and real workflows.

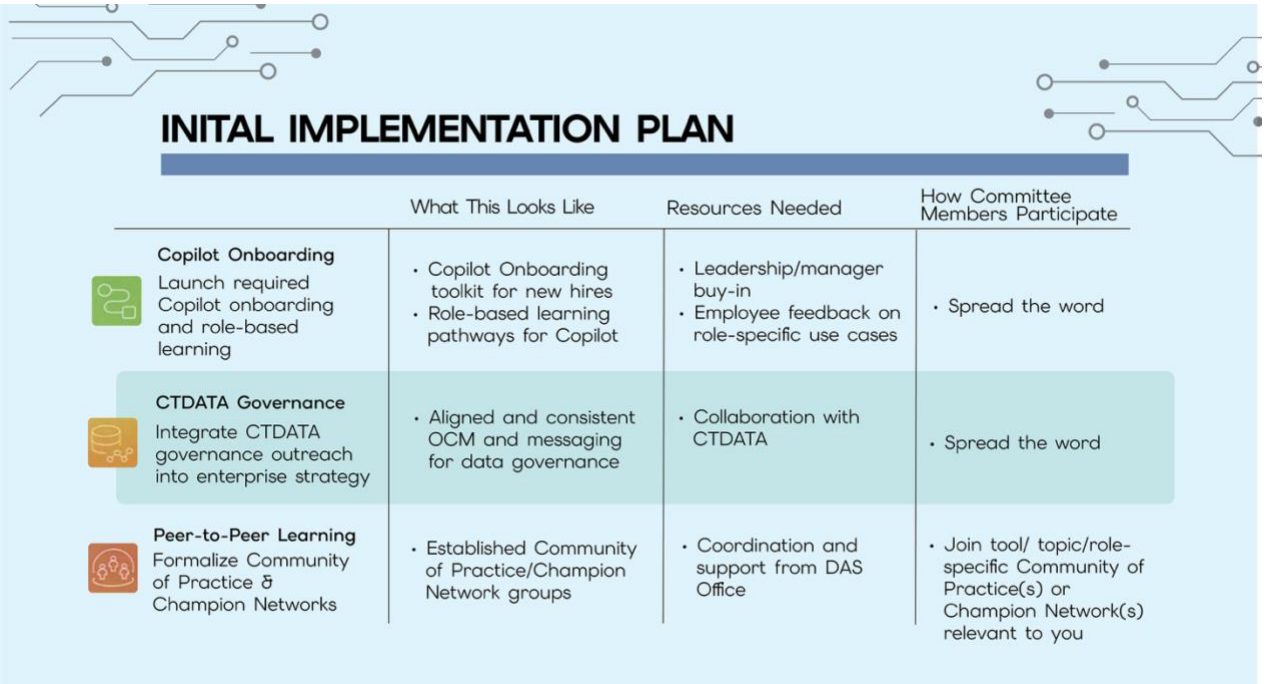

Figure 10. The Caltrans workforce-adoption roadmap: nine steps running from adoption of the U.S. Department of Labor AI literacy framework through required onboarding, communities of practice, and an AI center of excellence. Source: Caltrans presentation.

The most striking insight Shekhar offered was that a large part of the workforce had turned out to be far more advanced than expected, so much so that the program could not keep up with demand. He recounted a staff member who had built a working bridge-design assistant in two weeks and had to be asked to hold off on sharing it until it could be vetted, a small story that captured the central tension between grassroots innovation and responsible governance. His framing of adoption was that an agency should not force people through the door but open it and then show those who walk through what lies on the other side.

That framing was tested directly in the discussion. Asked whether AI use could fairly be made a performance expectation in a unionized environment, Shekhar argued that public-sector adoption should be tied to business outcomes rather than to the private-sector metrics of prompts per day or lines of code, and that the agency's job is to remove the real and perceived barriers to use while keeping a level playing field across job classifications. Unions, he added candidly, had been an early oversight on his part and were in fact a natural ally, since their interest in their members' capability aligns with the agency's; California law, he noted, was moving toward requiring thirty days' notice to unions before any AI feature is adopted.

The ITE AI Readiness Assessment

Ted Bailey presented the centerpiece of the session: a free self-assessment that lets an individual or an agency locate its current readiness and chart a path forward. He was careful about its purpose as the tool is educational rather than evaluative, collects no responses, and is built for transportation professionals. It runs as a live web tool reachable by a scan code and presents results on a maturity-model star diagram. ITE's AI in Transportation community developed the tool, drawing on the AASHTO TSMO capability maturity model, an established AI readiness framework, and a widely used commercial AI maturity model, adapted for the transportation context.

The assessment comes in two versions. The individual assessment poses fifteen questions across five dimensions: AI literacy, tool proficiency, transportation application, leadership and advocacy, and ethics and governance. It takes about six minutes to complete.

The agency assessment poses twenty questions across five dimensions: strategy and vision, data and infrastructure, workforce and culture, use cases and operations, and governance and ethics. This version takes about eight minutes to finish. Results map to named maturity levels, from an entry-level "explorer" profile to an "optimizing" profile for organizations that have made AI core to their operating model, each accompanied by suggested next steps. The room ran the assessment live and critiqued it usefully noting that the lower levels are described in fine gradations while the top levels feel monolithic, and suggesting the scale be flipped so that maturity becomes more granular as an agency advances. Bailey's framing of the tool's value was the line that stayed with the room: the value is not the score, but the conversation the score helps an agency have.

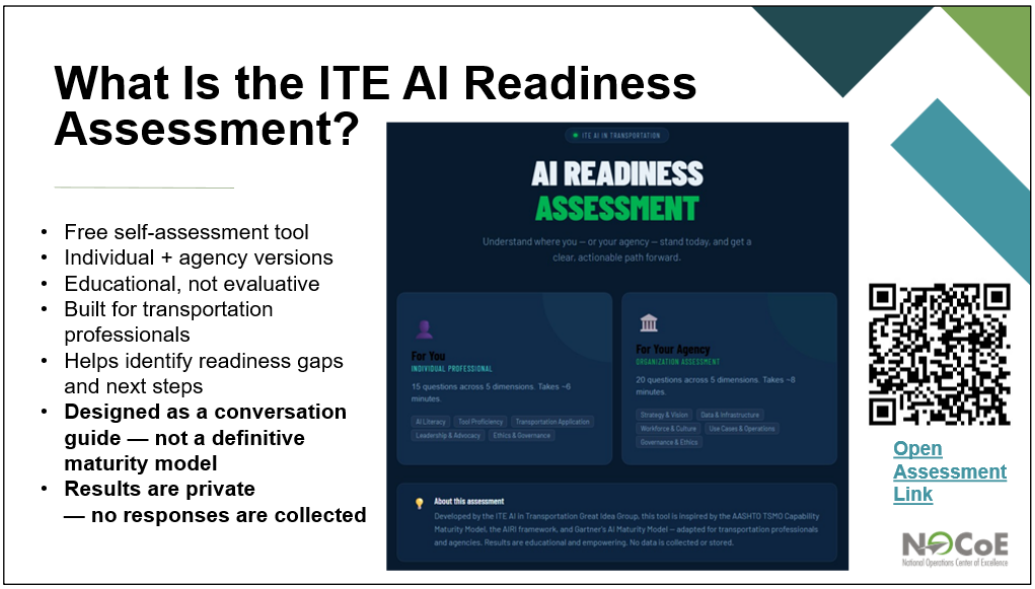

Figure 11. The ITE AI Readiness Assessment: a free, anonymous self-assessment in individual (15 questions, ~6 minutes) and agency (20 questions, ~8 minutes) versions, each scored across five dimensions. Participants took it live during the session. QR code in the figure provides link to open the assessment. Source: ITE / STV presentation.

Bailey also presented a role-based model of training, on the argument that AI readiness is not one-size-fits-all and that training should match a person's role, responsibility, and risk. The model runs as a tiered set of expectations: baseline AI literacy and acceptable-use awareness for all staff; safe day-to-day tool use for staff who work with AI directly; data and validation literacy for analysts, engineers, planners, and operations staff; oversight and governance for managers and supervisors; procurement and vendor-risk fluency for procurement, legal, and IT staff; and strategic risk and investment understanding for executives. The biggest gap, in his assessment, is precisely the role-based, transportation-specific training that would tell a DOT employee what to do differently depending on whether they are a new engineer, a traffic-management-center operator, a project manager, a procurement lead, a data steward, or an executive.

On training that agencies have actually delivered, the clearest model came from Texas DOT, which had offered, though not required, a Copilot training program to its first six hundred pilot users, running about eight weeks, one tool per week, beginning with the basics of prompt engineering and supported by a Teams channel for questions. Alongside that rollout, the agency has packaged a structured internal curriculum, the TxDOT AI Academy, summarized below. A separate cautionary thread concerned funding: an attempt to fund data work through a safety program was rejected as poor optics, with the lesson that data and AI need a dedicated funding source rather than a borrowed one, though participants noted that safety-specific data management can sometimes be funded through safety programs even when general data management cannot.

Inside the TxDOT AI Academy

The TxDOT AI Academy packages sixteen hours of coursework into four sessions, built largely from commercial learning-platform modules and sequenced deliberately from leadership framing to data discipline:

- **Session 1 — Orientation and guardrails**: Generative AI for Business Leaders; Prohibited Artificial Intelligence Practices; Demystifying AI; Streamlining Your Work with Microsoft Copilot.
- **Session 2 — The technical landscape**: Overview of Computer Vision; Introducing Robotic Process Automation; Artificial Intelligence Foundations: Machine Learning.
- **Session 3 — Generative AI in practice**: What Is Generative AI?; Your Top AI Questions Answered: AI Literacy for Everyone; Learning Microsoft 365 Copilot for Work; Ethics in the Age of Generative AI; Everyday AI Concepts.
- **Session 4 — Data culture**: Building a Data-Driven Culture; Data-Driven Decision-Making for the Real World; Global Data Stewardship; Learning Data Governance.

The arc of the curriculum is its quiet argument: it begins with what AI is and ends with data governance, the same data-first message the agency's practitioners delivered throughout the exchange.

McCulloch added a further production example from the procurement side of the house: an AI tool that scrapes data from invoices with hundreds of pages at a time into a database, cutting processing from hours to minutes. After roughly a year of use, it has reduced the late-payment fees the agency incurs, and the team's next step is matching invoices against contract rates to detect fraud, such as the same individual billing questionable hours across multiple contracts.

Roundtable on Key Learnings and Takeaways

Introduction to the Conversation

The exchange closed with each participant naming one key learning, one thing to take back to their agency, and one thing the group should keep working on together. The responses, taken together, form the clearest statement of what the two days produced, and the most important of them are synthesized here.

The recurring substantive theme was data. Doug Noble's reminder that the quality of any output depends on the quality of the input: garbage in, garbage out was echoed by AASHTO's call to treat data as an asset and to distinguish adoption from genuine capability, and by Texas DOT's insistence that data, like pavement and bridges, needs sustained funding to maintain. New Jersey DOT's contribution was that AI must be made predictable, and that data scientists should be brought onto pilots early, alongside the subject-matter experts, rather than after the fact. It also asked the leading states to share their workforce-training materials, which carry weight precisely because of where they come from. Michigan DOT pressed the importance of storytelling and breaking the work down into plain terms for the people who control budgets. Their own technology-branch strategic plan was in its final phase. Vermont's advice was to put data governance in place now, and to be wary of letting preoccupation with legacy systems foreclose future work; a second Vermont voice, eight months into the field, described positioning the state as an early-majority adopter rather than an innovator, content to learn from California and Texas. Tennessee DOT argued that AI, like TSMO before it, must be adopted by the organization as a formal practice even as individuals take initiative, and asked for recurring information on use cases on a faster cadence - every three or four months rather than a white paper a year after the fact. North Carolina DOT, inspired by the Caltrans safety tools, wanted to be a champion voice and pointed to a school-safety analysis that a similar tool could cut from a half-hour a week to a few minutes. Missouri DOT's word was exposure: to ideas, to questions, and to lessons learned, the good with the bad and a wish for continued conversation with NOCoE across the full lifecycle from concept through procurement to deployment. Joseph Gregory of FHWA saw opportunities to complement the TSMO capability maturity model with the federal AI work. One participant, capturing the room's recurring relief with a smile, called the exchange a professional therapy session.

Two unsolved problems were named explicitly and belong in any honest account of where the field stands: agencies do not yet have an accepted way to measure AI's effect on safety, or its effect on congestion and flow. The closing reactions captured the emotional register of the two days as much as the substantive one. Shekhar offered a perspective from the education and training profession about learning by doing: that one forgets what one is merely told, may remember what one is taught, but learns what one is involved in. This should be a frame, he argued, for how agencies should bring their workforces along. One observer noted what had been absent from two days of conversation: no one had asked for federal guidance to lead the way. The agencies were looking to each other.

When asked what NOCoE could most usefully do next, the room was consistent. It wanted more case studies, a compilation of the AI policies that states have already published, and more peer exchanges. Those requests shape the Next Steps that follow.

Final Thoughts

Kevin Riley closed for Caltrans, describing the two days as at once sobering, comforting, and disturbing. Sobering in the scale of the work ahead, comforting in the discovery that peers face the same unknowns, and disturbing in the pace at which the technology is moving. He cautioned that things may get worse before they get better in an age when everyone can claim, on the strength of a single article, to be an expert. His own contribution, he said, would be to pursue better-quality data and to be willing to say plainly when the data is not good enough to act on. He undertook to invite several of the agencies in the room, Virginia, Michigan, and Tennessee among them, along with partners from the consulting community, to brief his team.

Nicholas Ramfos closed for NOCoE with a set of commitments and a standing invitation. He pointed agencies to the resources already on the NOCoE website, including the prior peer exchange report on AI and computer vision, an AI strategy report, the recently updated Texas DOT strategic plan, the research on implementing machine learning at state DOTs, and NOCoE's library of model position descriptions, which agencies call asking for. He emphasized that NOCoE's role is as much to push as to pull, that he and his partners would get on a short screen-sharing call at any time to capture and share an agency's work, and that even a webinar with a handful of attendees can be a success when those few carry an idea home. With roughly forty webinars a year and open slots, the message to the room was simply to call when something works.

Key Learnings: an expanded view

The two days produced a set of cross-cutting themes that recurred across sessions and agencies. They are summarized here as the analytical core of the exchange.

Funding has become the binding constraint, and it is an operations problem as much as a startup problem. The leading concern moved from cost and over-reliance in 2025 to funding in 2026. The difficulty is not only securing money to begin but sustaining it: pilots can often be funded as capital, but the recurring costs of subscriptions, tokens, maintenance, and operations are hard to lock down, and formula funding does not reach them. Agencies want AI-specific funding categories in the US surface transportation reauthorization legislation rather than the general language of "digital infrastructure," and several are experimenting with their own answers, from Vermont's low-cost in-house model to per-project token chargebacks.

Data is the precondition, and it should be treated and funded as an asset. Across Caltrans, Texas DOT, AASHTO, and New Jersey DOT, the same point recurred: there is no successful AI without a sound data foundation. The most concrete lesson of the exchange came from Texas DOT in that free text is the enemy of usable data, that an agency may carry hundreds of variations of a single road name, and that the unglamorous work of standardizing data and joining incompatible systems is what makes everything downstream possible. The AASHTO survey put numbers to the same point: data quality and availability is the single most-cited challenge in the field. Treating data as an asset means standardizing inputs at the source, governing data deliberately, scanning out sensitive information before models touch it, verifying what models are fed, and funding the upkeep of data the way an agency funds the upkeep of physical infrastructure.

Workforce and organizational readiness are the binding human constraint and the questions run deeper than training. The scarce profile is the person who can bridge civil and traffic engineering with coding, data, and AI. Agencies are addressing the gap through readiness assessments, role-based training, productivity-tool cohorts such as Texas DOT's, and tools that capture the knowledge of retiring staff, but capability is fragile, as Texas DOT's reduction from two people to one made this point. The deeper obstacles are institutional: human-resources classification systems that misread new skills, the role of unions as both stakeholder and ally, and a talent pipeline that demands experience for jobs that are supposed to provide it. The field still lacks role-based, transportation-specific training, and a surprising amount of capability already exists, untapped, in the current workforce.

Governance turns on the distinction between decision support and automated decision-making. The workable standard the group converged on is that AI should support a knowledgeable human's judgment rather than make the decision, that the review which follows must be genuine. Caltrans requires sign-off from a substantial majority of its safety engineers and that accountability, including for contractors acting on an agency's behalf, survives the introduction of AI. The assertion that a model made the decision will not satisfy a court, and documenting an AI-surfaced hazard creates its own duty to act.

Contracting is a current and real, unresolved problem. AI hides inside tools agencies already buy; most agencies lack a working definition of AI; and policies that address only generative AI leave specialized machine learning unaddressed. The active fronts are vendor vetting and performance testing, the prohibition on sole-source procurement, the protection of agency data and intellectual property when models are trained on agency data, the two-tier accountability gap introduced by systems integrators, and the allocation of liability.

The supporting ecosystem is maturing into the clearinghouse agencies are asking for. Agencies were candid about not knowing what they did not know, and about their exposure to vendor claims they could not evaluate. The partners are beginning to answer that: AASHTO's survey, white paper, dashboard, and forthcoming maturity framework; the federal AI Center of Excellence and its initiatives on intersection safety and traffic management centers; ITS America's implementation guide, governance guides, self-assessment, and use-case library; ITE's readiness assessment; and the multi-state pooled funds. The pieces of a national clearinghouse exist; the work ahead is to connect them.

Multi-state models lower the barrier to entry. Pooled funds such as the Michigan-managed ENTERPRISE program and the New Jersey-led State DOT Council for Strategic AI Adoption distribute financial risk and create shared frameworks and a common language, an advantage that matters most to smaller agencies, and the State DOT Council offers a ready home for sharing the use-case inventories that larger agencies have already built.

Demonstrating value is itself a capability the field has not mastered. Agencies named themselves, candidly, as poor storytellers. They recognized that translating AI's benefits into terms that resonate with leadership and the public is as much a part of the work as the technology. This can include producing the case studies that show savings, as Texas DOT did with a multimillion-dollar efficiency estimate, as well as a solid business case that shows the impact to customers.

Two measurement problems remain open. The field has no accepted method for measuring AI's effect on safety or its effect on congestion and flow. Until it does, the value case will rest on narrative and proxy rather than on demonstrated outcomes.

The posture of the field is bottom-up and agency-led. Perhaps the most distinctive feature of the exchange was the absence of any call for federal direction. Agencies looked to one another, and to peer models such as those in Texas and Utah, with the federal partners offering resources and coordination rather than a mandate.

Funding, data, and workforce readiness are the core constraints on AI adoption. Human oversight remains essential, while contracting, liability, and sustainable operating models are still evolving. Peer learning, pooled investments, and shared resources are helping agencies move forward, but the field still needs stronger ways to measure and communicate AI's value. Notably, progress remains largely bottom-up and agency-led rather than federally directed.

Next Steps

NOCoE intends to brief AASHTO, ITE, ITS America, and FHWA on the findings of this exchange and to pursue the products and follow-on work the participants requested. The next steps below are organized around NOCoE's convening role, the products it can produce, the partner efforts agencies can join, and the specific commitments participants offered.

NOCoE's convening and clearinghouse role. Participants were consistent in wanting NOCoE to continue the AI for TSMO conversation and, in particular, to stand up or point to a shared library of use cases, the single most-requested resource, and the direct answer to the vendor anxiety agencies described. NOCoE will also provide recurring updates on use cases on a faster cadence than an annual report allows, in response to the explicit request for information every three to four months and will refresh its model TSMO position descriptions for the AI era, including roles such as the artificial intelligence scientist, the transportation data ethicist, and the data management specialist, with input drawn through the Technical Advisory Committee and ITE's practitioner community. NOCoE also reiterated its standing offer to capture and broadcast an agency's work through a short screen-sharing session or one of its roughly forty annual webinars.

Products NOCoE can produce. Several products map directly to participant requests: a compilation of the AI policies that states have already published, so that an agency drafting its own need not start from a blank page; a series of case studies that demonstrate cost savings and operational value, written for the leaders and the public who must be persuaded; a contracting playbook for AI in TSMO, covering definitions, enforceable requirement language, data-ownership and intellectual-property clauses, liability, vendor scoring, and the no-sole-source pathways agencies are using; and a primer on treating data as an asset, paired with guidance on data literacy and the role-based, transportation-specific training the field still lacks. A risk-acceptance framework, a gap Texas DOT and AASHTO both named explicitly, would give agencies a structure to adapt to their own codes and policies. Three actions were recorded on the spot during the closing discussion: a technical brief on data governance; a multi-segment webinar series on basic AI tools, developed with ITE and the partner associations; and additional peer exchanges that carry the AI conversation beyond TSMO into planning and traffic engineering.

Partner efforts agencies can join. Several existing vehicles offer agencies a way to go further together. AASHTO maintains its national survey, white paper, and dashboard and is a candidate home for the best-practice study on ITS and TSMO organization that participants requested. The federal AI Center of Excellence offers a national clearinghouse, the intersection-safety datasets on the federal data hub, and the traffic-management-center benchmark work, and is seeking state case studies. ITS America offers an AI implementation guide, governance group guides, a self-assessment worksheet, a use-case library, a forthcoming privacy guide, and an academy course. ITE offers its AI readiness assessment and its practitioner community. Agencies seeking to coordinate across state lines can join the Michigan-managed ENTERPRISE pooled fund or the New Jersey-led State DOT Council for Strategic AI Adoption, whose first in-person peer exchange is planned for Dallas, Texas.

Specific commitments and offers. Caltrans offered to share use cases that might spark ideas elsewhere, pointed to its public AI program site, and committed to bringing several peer agencies in to brief its team. Texas DOT offered its emerging-technology pilot program as a model for vendor vetting and its Copilot training cohort and AI Academy curriculum as models for workforce development. Iowa DOT offered its candid pilot post-mortems as cautionary case studies. ITS America offered its full set of resources explicitly for inclusion in this report. Vermont offered its custom in-house model as a low-cost approach worth a closer look. Virginia DOT asked the ITE Education Council to help connect data-science coursework and student chapters to the agency talent pipeline and pointed participants to active research on the topic through the Transportation Research Board.

The next phase is to turn shared learning into practical tools agencies can use. NOCoE can play a central convening and clearinghouse role by expanding peer exchange, maintaining a shared use-case library, and developing resources on contracting, data governance, workforce readiness, risk, and demonstrating value. At the same time, agencies can build on existing partner initiatives and multi-state collaborations rather than starting from scratch. The emphasis is on connecting what already exists, filling the most persistent gaps, and making proven approaches easier to find, adapt, and scale.

Resources

The following publicly available resources were shared or referenced during the exchange.

Caltrans
Caltrans Artificial Intelligence program: https://innovation.onramp.dot.ca.gov/artificial-intelligence-caltrans

US DOT / FHWA
FHWA AI Center of Excellence: https://highway.dot.gov/ai-coe

ITS DataHub (including intersection-safety training and validation data): https://www.its.dot.gov/data/

Multi-state and pooled-fund efforts
ENTERPRISE Pooled Fund — *Role of Artificial Intelligence (AI) in Intelligent Transportation Systems (ITS)*: https://enterprise.prog.org

Transportation Research Board
Active research referenced by Virginia DOT:
https://apps.trb.org/cmsfeed/TRBNetProjectDisplay.asp?ProjectID=5704

ITS America
A Guide to Practical Next Steps for AI Implementation: https://itsa.org/wp-content/uploads/2025/03/2025_ITSA_AI-Implementation_GuideFinal.pdf

AI Self-Assessment Worksheet (Path to AI Maturity): https://itsa.org/wp-content/uploads/2026/05/AI-Self-Assessment-Worksheet.pdf

AI Governance Group Guides: https://itsa.org/advocacy-material/ai-governance-group-guides/ Technology Use Case Library: https://itsa.org/technology-use-case-library/

ITS America Academy — Artificial Intelligence in Transportation (course): https://itsa.org/wp-content/uploads/2025/10/AI-Course-Flyer.pdf

NOCoE
AI for TSMO Peer Exchange Report (2025) and other peer exchange reports:
https://transportationops.org/reports

Model TSMO Position Descriptions: https://transportationops.org/workforce/model-tsmo-position-descriptions

Other
Texas DOT Artificial Intelligence Strategic Plan: https://www.txdot.gov/content/dam/docs/division/str/ai-strategic-plan-09-20-2024.pdf

NCHRP research on implementing machine learning at state DOTs:
https://nap.nationalacademies.org/catalog/27880

Appendix A — Agenda

Artificial Intelligence for TSMO — May 19–20, 2026 — Sacramento, CA

Day 1 — Tuesday, May 19

Topic	Time	Speakers
NOCoE Welcome; why are we having this peer exchange; agenda approach and ground rules	9:00 – 9:15 a.m.	Nicholas Ramfos, NOCoE
Host Welcome and Caltrans AI Vision	9:15 – 9:45 a.m.	Kevin Riley, Caltrans; Vidhu Shekhar, Caltrans
Caltrans AI Projects: Traffic Mobility Insights; Road Safety Insight Platform; research projects	9:45 – 10:45 a.m.	Rajpreet Khangura, Caltrans; Jaswinder Gill, Caltrans; Prakash Sah, Caltrans
Break	10:45 – 11:00 a.m.	
AI in Pooled Fund Studies: Enterprise AI in Transportation; State DOT Council on Strategic AI Adoption	11:00 a.m. – 12:00 p.m.	Marlon Spinks, Michigan DOT; Giri Venkiteela, New Jersey DOT
Lunch	12:00 – 1:00 p.m.	
Partner Initiatives: AASHTO AI studies; ITE Great Idea Group; ITS America Committee; US DOT	1:00 – 2:30 p.m.	Ahnaf Morshed, AASHTO; Ted Bailey, STV (NOCoE TAC); Charnita Allen, DDOT; Kevin Viita, ITS America; Joseph Gregory, US DOT
Break	2:30 – 2:45 p.m.	
DOT Roundtable on Uses of AI	2:45 – 4:00 p.m.	All attendees
Key Learnings and Preview of Tomorrow	4:00 – 4:30 p.m.	Nicholas Ramfos, NOCoE

Day 2 — Wednesday, May 20

Topic	Time	Speakers
NOCoE Welcome; recap of Day 1; goals for today	9:00 – 9:15 a.m.	Nicholas Ramfos, NOCoE
Expanding Use of AI in Your DOT: leadership engagement; business case; leveraging existing programs	9:15 – 10:45 a.m.	Eryn Stone and Matthew Miller, Iowa DOT; Ben McCulloch, Texas DOT
Break	10:45 – 11:00 a.m.	
AI in Contracting: best practices; data analysis; agency coordination	11:00 a.m. – 12:00 p.m.	All attendees
Lunch	12:00 – 1:00 p.m.	
Workforce and Training; data literacy	1:00 – 2:30 p.m.	Ted Bailey, STV (NOCoE TAC); Vidhu Shekhar, Caltrans
Break	2:30 – 2:45 p.m.	
Roundtable on Key Learnings and Takeaways	2:45 – 4:00 p.m.	All attendees
Final Thoughts and Next Steps	4:00 – 4:30 p.m.	Kevin Riley, Caltrans; Nicholas Ramfos, NOCoE

Appendix B — Participants

The exchange convened practitioners from state departments of transportation, the United States Department of Transportation and the Federal Highway Administration, the NOCoE Technical Advisory Committee, and partner associations. The participants below are drawn from the session records.

Name	Organization
Kevin Riley	Caltrans (host)
Vidhu Shekhar	Caltrans
Rajpreet Khangura	Caltrans
Jaswinder “Sunny” Gill	Caltrans
Prakash Sah	Caltrans
Eryn Stone	Iowa DOT
Ben McCulloch	Texas DOT
Marlon Spinks	Michigan DOT
Jason Bodell	Michigan DOT
Giri Venkiteela	New Jersey DOT
Sanhita Lahini	Virginia DOT
Rose Vath	Vermont Agency of Transportation
Greg Dyer	Tennessee DOT
Seth Daniels	Nevada DOT
Aiden Burnside	North Carolina DOT
Rick Zygowicz	Missouri DOT
Stin Weber	Maricopa County (Arizona)
David Man	Metropolitan Transportation Commission
Henry Hammel	Metropolitan Transportation Commission
Sarah Burnworth	Metropolitan Transportation Commission
Charnita Allen	District Department of Transportation
Joseph Gregory	Federal Highway Administration
Govind Vadakpat	Federal Highway Administration
Ahnaf Morshed	AASHTO
Ted Bailey	STV (NOCoE Technical Advisory Committee)
Kevin Viita	ITS America
Douglas Noble	ITE
Nicholas Ramfos	NOCoE
Adam Hopps	NOCoE