

April 2026

RESEARCH PROJECT TITLE

Improving Work Zone Management and Safety through AI-Powered Connected Vehicle Data Analysis

SPONSORS

Smart Work Zone Deployment Initiative
(Part of TPF-5(438))
Federal Highway Administration
(InTrans Project 24-888)

PRINCIPAL INVESTIGATOR

Anuj Sharma, Co-Director
Real-Time Analytics of Transportation Data (REACTOR) Laboratory
Iowa State University
anuj@iastate.edu / 515-294-3624
(orcid.org/0000-0001-5929-5120)

CO-PRINCIPAL INVESTIGATOR

Skylar Knickerbocker, Research Engineer
Center for Transportation Research and Education, Iowa State University
(orcid.org/0000-0002-0202-5872)

MORE INFORMATION

intrans.iastate.edu

SWZDI
Iowa State University
2711 S. Loop Drive, Suite 4700
Ames, IA 50010-8664
515-294-8103

The Smart Work Zone Deployment Initiative (SWZDI) is a transportation pooled fund that supports research investigations into better ways to improve the safety and efficiency of traffic operations and highway work in work zones. The primary objective is to promote and support research and outreach activities that focus on innovative policies, processes, tools, and products that enhance the implementation, safety, and mobility impacts of work zones. The fund is administered by Iowa State University's Institute for Transportation, and the lead agency is the Iowa Department of Transportation.

The sponsors of this research are not responsible for the accuracy of the information presented herein. The conclusions expressed in this publication are not necessarily those of the sponsors.

Improving Work Zone Management and Safety through AI-Powered Connected Vehicle Data Analysis

tech transfer summary

ReactorIQ is an integrated, cloud-based platform that leverages connected vehicle data to enable real-time incident detection, daily operations analysis, and data-driven support for smart work zone planning.

Objective

This project set out to design and implement a comprehensive statewide, real-time work zone intelligence system, ReactorIQ, capable of transforming raw connected vehicle (CV) telemetry, camera feeds, and cloud-based analytics into actionable operational insights.

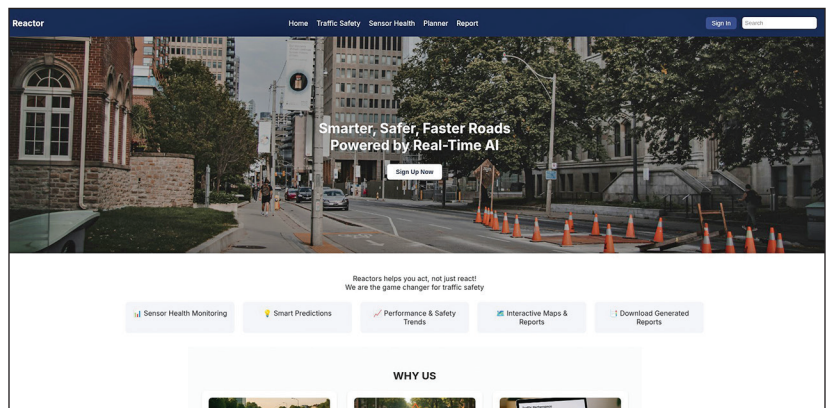
Background and Problem Statement

Work zones can significantly disrupt normal traffic flow through, for example, lane shifts, narrowed shoulders, and tapered lane drops. If not monitored and managed carefully, these disruptions can escalate into safety and mobility problems.

Traditional fixed and temporary monitoring systems, including loops, cameras, and radar units, are often insufficient by themselves to monitor congestion and queues. These systems lack complete spatial coverage and have limited ability to capture emerging conditions across rural corridors and secondary roadways.

The expansion of CV technologies in recent years offers a promising solution for work zone monitoring and management. Advanced analytical tools can leverage the large volumes of CV telemetry data generated by passenger vehicles and commercial fleets to capture queue dynamics and bottleneck formation in real-world conditions.

These developments provide a foundation for integrated platforms that combine CV data with traditional sensors, probe products, and emerging field devices to support more robust and proactive work zone management.



ReactorIQ dashboard landing page

Description of ReactorIQ

ReactorIQ is based on a cloud-native analytics architecture capable of processing millions of CV records per hour with near-real-time latency. Built on AWS Glue, Apache Spark, and Apache Sedona, the system filters and maps telemetry onto roadway geometries and generates several performance indicators.

The system is designed around three core pillars: real-time incident detection, daily operations analysis, and data-driven support for smart work zone planning.

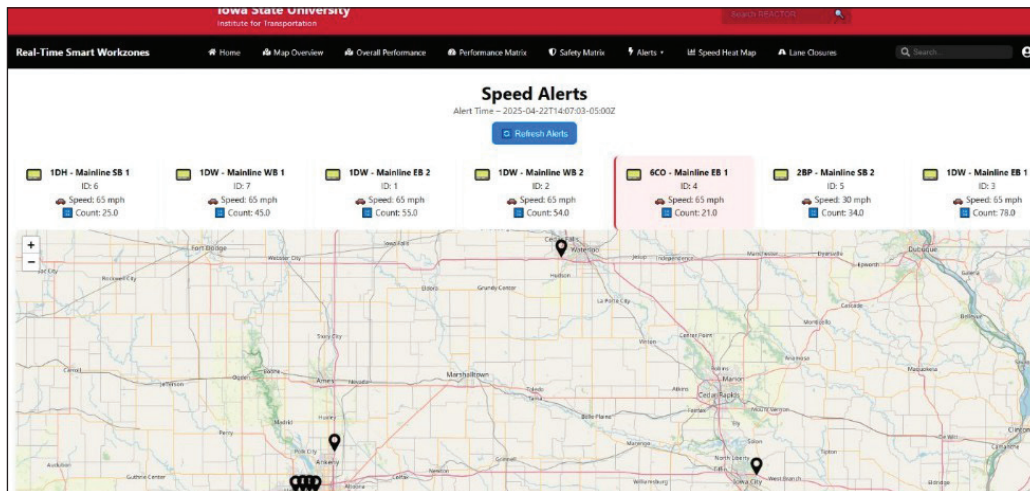
Real-time incident detection involves revealing emerging slowdowns, queue buildups, speed drops, and turbulence patterns on the order of minutes. CV telemetry data are ingested into cloud services that detect conditions associated with elevated crash risk, and messages are generated for use in dynamic message signs (DMS) and other applications.

Daily operations analysis provides the operational layer of the system. ReactorIQ continuously computes CV-based performance metrics and presents them through a web-based dashboard that visualizes statewide work

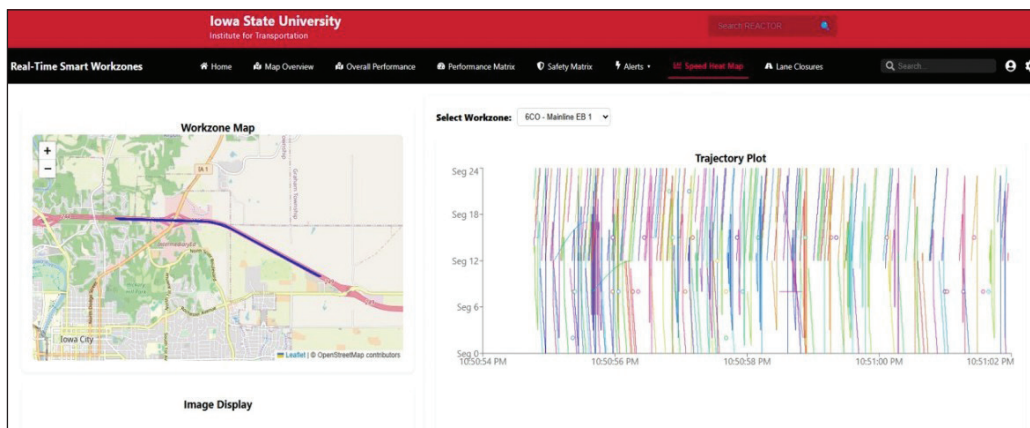
zones, real-time speed segment heatmaps, daily time series plots, and vehicle trajectory diagrams.

For smart work zone planning, historical CV telemetry data and hybrid modeling frameworks are leveraged to predict how flow, density, speed, and queue length evolve under baseline and work zone conditions. The outputs are used to examine alternative deployment windows, lane closure configurations, and staging strategies. The dashboard also incorporates outputs from a CV-based automatic lane closure detection model, which infers closure status from vehicle behavior.

Collectively, these three pillars establish a CV-driven framework that elevates ReactorIQ from a monitoring dashboard to an artificial intelligence (AI)-driven work zone management system. The resulting system provides transportation agencies with powerful tools for detecting congestion, accurately identifying lane closures, monitoring safety-related behavior, and understanding traffic dynamics at both the macroscopic and microscopic levels.



Real-time speed alerts in ReactorIQ



Real-time vehicle trajectories within a work zone in ReactorIQ

Key Results

Challenges

- An early challenge was to design robust spatial queries capable of accurately joining massive volumes of CV telemetry to the underlying roadway network. Early versions of the pipeline frequently encountered spatial mismatches, leading to incorrect assignment of waypoints to segments.
- In low-volume rural corridors with sparse CV penetration, naive algorithms that rely on high sample densities or strict thresholds on mean speed can underestimate queues and slowdowns or produce fragmented detections. This was mitigated by using percentile-based speed metrics and temporal persistence rather than raw averages alone. This allowed the platform to detect meaningful slowdowns while avoiding overreaction to isolated slow vehicles.
- Early approaches to lane closure detection yielded spurious flags in locations where traffic was normal but sample sizes were small. Reliability was improved substantially through a more robust representation based on empirical cumulative distribution functions (CDFs), Kolmogorov-Smirnov statistics, and multi-week aggregation.
- Early prototypes featured high latency and unsustainable costs for long-term, always-on operations. Transitioning the majority of heavy processing to AWS EC2-based Spark clusters reduced latency, considerably lowered monthly compute costs, and provided finer control over scaling during peak and low-demand periods.
- Ensuring end-to-end reliability of the real-time pipeline required operational refinements. Some failures arose when services attempted to consume files that had already been deleted or overwritten. To resolve this, dual-write strategies were adopted, retention windows were extended, and pipeline stages were synchronized using event triggers.
- Dashboard responsiveness and map rendering performance initially suffered when thousands of CV points and polygons were drawn at full resolution. This was mitigated through map tile caching, asynchronous data loading, and zoom-dependent filtering so that the dashboard displays only the level of detail appropriate for the current view.

Accomplishments

- The fully automated real-time pipeline is capable of ingesting, processing, and analyzing CV data on a continuous basis. The system routinely processes CV telemetry on the order of every 30 seconds and generates statewide traffic performance metrics that capture queue formation, lane closures, and sudden speed drops with minimal latency.

- A unified operational dashboard integrates several analytics into a cohesive interface for transportation operators. Users can view statewide work zones on an interactive map, inspect real-time speed heatmaps and queue metrics, review time series plots for speed and delay, and examine high-resolution trajectory plots of microscopic driving behavior. The dashboard further exposes lane closure detection overlays and, when available, live camera feeds for visual verification.
- A particularly impactful success was the accurate detection of lane closures using only CV data. By continuously analyzing CV lateral position and speed distributions at a 0.1-mile resolution, ReactorIQ provides an independent behavior-based source of truth for actual closures. This improves day-to-day situational awareness and offers an accurate basis for evaluating work zone performance and contractor compliance over time.
- The system generates real-time queue warnings and safety alerts through XML messages describing the onset and extent of queues and severe slowdowns near work zones. Compliant with Traffic Management Data Dictionary (TMDD) standards, these messages are designed to feed directly into agency systems for DMS activation and traveler information dissemination, supporting active, automated responses to emerging traffic conditions.
- The diagnostic value of visual trajectory plots and speed heatmaps for operational decision-making was demonstrated. These tools have already provided insights into the precise characteristics of breakdowns, queues, and recoveries and directly support both immediate operational decisions and longer-term work zone design improvements.
- The architecture itself represents a major success. By shifting heavy computation to AWS EC2-based Spark clusters and designing modular services, the system achieves statewide scalability while maintaining low latency and reasonable operating costs. The modular design also makes it feasible to extend ReactorIQ to additional use cases.

Future Development

Future work can extend the platform in several meaningful directions:

- CV telemetry data can be fused more deeply with camera-based analytics. AI-based video models capable of identifying stalled vehicles, estimating lane occupancy, and detecting near-miss interactions could significantly improve the detection of complex incidents and subtle safety issues, particularly in areas with lower penetration rates.

- The lane closure detection framework can be extended to handle dynamically moving closures, partial shoulder closures, multi-stage construction layouts, and highly variable nighttime operations.
- Improvements to refine trajectory smoothing, enhance shockwave detection algorithms, and formalize turbulence classification into clear operational categories would help operators quickly interpret complex patterns and prioritize responses.
- Work is planned to strengthen data redundancy, incorporate message queues for more resilient inter-service communication, and further optimize storage and compute costs. Additional use of AWS EC2 autoscaling policies and spot instances may provide further cost savings without compromising reliability.
- Several user-facing enhancements are envisioned for the dashboard. These include views for comparing multiple work zones across corridors, historical playback modes for replaying past days of operation, and automated periodic performance reports that summarize key metrics and anomalies.

Implementation Readiness and Benefits

The project's technical and operational accomplishments demonstrate the feasibility and value of a statewide work zone intelligence system based on CV data, cloud-native processing, and tailored analytics. The successful deployment of ReactorIQ represents a significant advancement in real-time traffic monitoring and work zone management

With continued enhancements, ReactorIQ can serve as a foundation for statewide automated operations, predictive analytics, and integrated work zone management, laying the groundwork for the next generation of smart transportation infrastructure in Iowa and beyond.