PRESENTED TO THE FHWA DIGITAL DELIVERY STAKEHOLDER GROUP — MEETING #4, JULY 2026

"How Edge Compute Enables Real-Time Applications and Supports Agency Data Management."

The Data Beneath the Road

Vikasa · OpenITS — An Open Data Fabric for Transportation Infrastructure.

Michelle Maggiore

2M Transformation Services for FHWA Digital Infrastructure Enablers (HRDI-ADC-11)

WHY THIS MATTERS RIGHT NOW

Open models won. Open roads are next.

On July 24, 2026, NVIDIA and a coalition of dozens of AI companies — Microsoft, Meta, IBM, GitHub, the Linux Foundation, Perplexity, and more — published "Open Weights and American AI Leadership," arguing that open, inspectable models beat closed ones on access, competition, and even safety: openness lets many teams find and fix flaws, while closed systems concentrate risk in a few hands.

> *"Open weights expand access to the AI economy."*

> *"Openness may be one of the most important paths to AI safety and security."*

> *"Concentrating advanced capabilities behind a small number of closed models compounds risk."*

Transportation data infrastructure faces the identical choice — open, industry-governed schemas versus closed, single-vendor lock-in. The AI industry just made its case publicly. This deck makes the same case for the data beneath the road.

Source: NVIDIA, "Open Weights and American AI Leadership" (July 24, 2026)

WHO THIS AFFECTS

Lifecycle data, not just live traffic.

DDSG's own charter calls for "**lifecycle data management that support semantically rich digital twins**" — this is that layer.

Inside a DOT

Traffic operations, asset & BIM management, IT/GIS, planning, and digital delivery & data governance teams — each currently locked into different vendor schemas for the same physical assets.

The traveling public

Faster incident response, fewer crashes ($340B annual crash cost, 42,795 deaths — NHTSA), more reliable travel times, and transparency into how public infrastructure data is used — outcomes only possible when the underlying data can move freely across systems.

A canonical, open schema for real-time and lifecycle data is the foundation both a digital twin and a traffic signal read from — this group sits at exactly that junction.

vikasa

THE FEDERAL GROUNDWORK BEHIND THIS ROOM

FHWA ran the research. VikasaOpenITS is what it pointed to.

"How Edge Compute Enables Real-Time Applications and Supports Agency Data Management."

30+

vendor conversations across signal, sensor, and edge-compute suppliers, run under USDOT's Digital Infrastructure Enablers program

63 use cases

mapped across traffic optimization, planning, freight, security, digital twins, and asset management — the market, not a hypothesis

9 layers

the reference architecture that came out of it — field devices to application, with one governed data plane in the middle

Four practitioner workshops, six tech memos, and a national codefest converged on one finding: data interoperability — not another device standard — is the layer transportation is missing. Vikasa is that data plane, built and running.

Source: 2M Transformation Services, “Why Digital Infrastructure Needs a Shared Data Layer” — FHWA Digital Infrastructure Enablers (HRDI-ADC-11), March 2026

THE UNCOMFORTABLE TRUTH

You own the roads.
You do not own the data.

Every agency runs the same classes of roadside infrastructure — signal controllers, dynamic message signs, environmental sensors, roadside V2X units, ramp meters. Yet the telemetry those devices produce is trapped inside each vendor's silo, invisible across agency and jurisdiction boundaries.

Locked in

Data lives in proprietary formats controlled by whoever sold the box — not by the agency that paid for it.

Locked out

Neighboring agencies on the same corridor cannot see each other's real-time signal state or incidents.

Locked down

Every new integration is a custom, months-long project. Innovation waits on procurement.

vikasa

THE STAKES ARE MEASURED IN LIVES AND DOLLARS

The status quo has a price tag.

$340B

annual economic cost of U.S. motor-vehicle crashes (2019)

$1.4T

total societal harm when quality-of-life losses are included (2019)

42,795

people killed on U.S. roads in a single year (2022)

7,314

pedestrians killed — the highest count in four decades (2023)

These numbers do not move because the systems that could prevent them cannot talk to each other. Safety is an integration problem before it is an engineering problem.

Source: NHTSA — Economic & Societal Impact of Crashes, NHTSA FARS

THE INDUSTRY ALREADY AGREES ON THE PROBLEM

Four challenges block modern DOT data collection and sharing. Three have solution paths. One does not.

There are four hard problems facing intelligent transportation networks. The market has products for real-time processing, security, and scale. Interoperability has been left as an afterthought — and it is the one that blocks everything else.

ADDRESSED

Real-Time Processing

Edge compute & AI models

ADDRESSED

Security & Privacy

Firewalls, SIEM, zero-trust

ADDRESSED

Scalability

Cloud, industrial networking

THE GAP WE FILL

Interoperability

Protocol & standards alignment

vikasa

THE ONE ARCHITECTURAL IDEA THAT CHANGES EVERYTHING

Decouple the producers of data from the consumers of data.

Instead of wiring devices directly to applications, normalize every device's output into shared, canonical event types published over a common messaging layer. Producers speak one open language. Consumers subscribe to meaning, not to hardware.

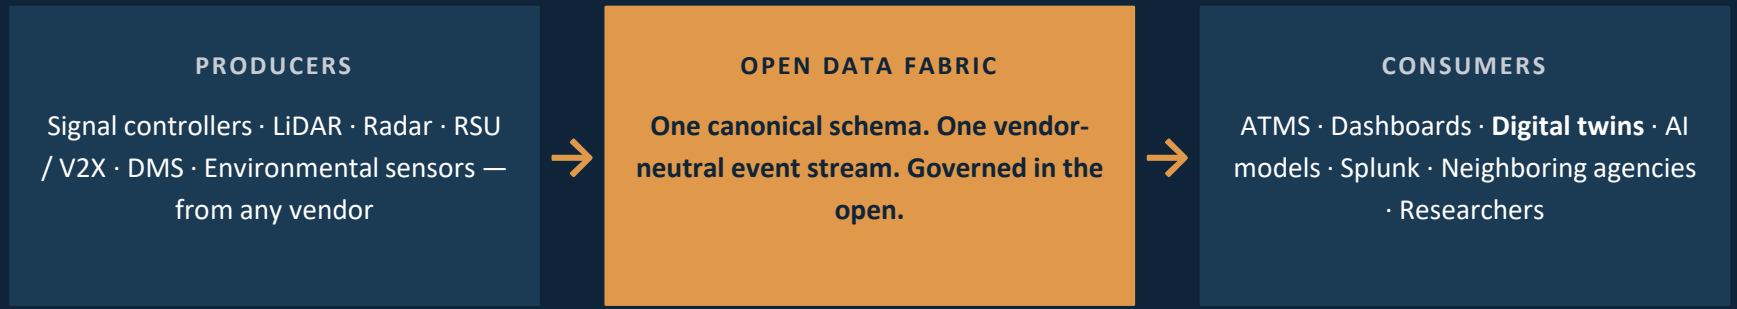

Digital twins subscribe to the same canonical stream as ATMS and analytics — not a separate export.

THE STANDARDS WERE BUILT FOR A WORLD THAT NO LONGER EXISTS

Transportation runs on 30-year-old plumbing.

NTCIP, SAE J2735, and ARC-IT were all conceived between 1992 and 1996 — and NTCIP still rides on SNMP/SMIv2, whose data-modeling roots trace to the 1988 version of ASN.1. We are carrying connected-vehicle, digital-twin, and AI-at-the-edge workloads on a foundation designed before the modern web existed.

1988
ASN.1 — the encoding NTCIP still inherits

1992
NTCIP work begins (NEMA)

1996
National ITS Architecture (now ARC-IT) published

2010
YANG published RFC 6020

2-3 yrs
to adopt a single new function — NTCIP 1202 v04 ran Aug 2022 to Aug 2025

5-10 yrs
standards-review cadence at SAE (5) and IEEE (10) before a refresh is even required

30 yrs
since the ITS data foundations were designed — patched forward, never rebuilt

Source: Wikipedia (NTCIP), ITE NTCIP Protocol Task Force (1202 v04), SAE Standards Definitions, IEEE SA, ARC-IT

vikasa

OUR CONTRIBUTION

Vikasa is that open data fabric — built, not theorized.

Vikasa is two things in one project, because a specification you cannot run is a specification you cannot trust.

An open standard

- A YANG canonical data model
- CloudEvents 1.0 envelope + Protobuf payloads
- A 7-token open subject hierarchy (OpenITS)
- Governance rules + an extension framework

A runnable reference implementation

- An edge poller for roadside cabinets
- A central ingest service (ClickHouse / KV / Prometheus)
- A conformance harness that tests compliance
- Tooling to add new device classes without forking

Aligned with NTCIP, SAE J2735, IEEE 1609, and ARC-IT — filling the gaps those standards leave open (no registry, no schema versioning, no graduation path).

vikasa

THIS IS NOT A WHITEBOARD CONCEPT

At the April 2026 FHWA CodeFest, it worked — in days, not months.

Days

to stand up a multi-vendor canonical data fabric — not the months a proprietary integration would take

40+ stakeholders

across agencies and vendors, connected to the same shared fabric and executing end-to-end actuation through canonical schemas

9 layers

a validated edge-to-cloud reference architecture: field devices → pollers → cabinet nodes → fabric → apps

Real adapters, safety logic, and live edge-to-cloud actuation on real data streams. The output was not a slide deck — it is working schemas, code, runbooks, and governance patterns any agency can reuse.

Source: FHWA CodeFest Playbook (Deliverable D26) — 2M Transformation Services & 360ns

A DEMONSTRATION OF THE PATTERN

Sovereignty by default. Sharing is opt-in.

An open fabric does not mean surrendering control. Each agency runs its own stack — cabinet edge, regional cluster, org core, optional DMZ. Cross-agency sharing happens only through explicit, subject-filtered peer subscriptions. No central broker. No shared database. No agency cedes control of its own infrastructure.

I-85 Multi-State Data Corridor Lab

GDOT, SCDOT, and NCDOT federating across state lines on the same interstate corridor — sharing real-time signal status without a central broker and without any agency ceding control of its network.

This was a synthetic-data lab exercise proving the federation architecture works — not a live production deployment.

3 DOTs
GDOT · SCDOT · NCDOT

23 data sources
federated across the lab

9 locations
on the shared I-85 corridor

0
shared databases or brokers

Source: Vikasa architecture docs — federation & corridor lab

vikasa

SEE IT RUNNING — LIVE, ON OPEN SOURCE, ON YOUR DATA

One open corridor. Five things no ATMS can do today.

WHAT YOU'RE SEEING

1. **AI reads your data because it's modeled**
   Given only our OpenITS YANG model, an AI agent explores the fabric and builds a working dashboard. Open source, no proprietary schema.
2. **GDOT managed lanes, finally visible**
   I-75 reversible express lanes stream live open-direction state — the view the ATMS can't give them.
3. **ATSPM from raw controller events**
   Purdue diagrams and split-failure analytics derived on the fly — no vendor ATSPM module in the loop.
4. **3 states, one live corridor**
   Camera + lidar fused across GA/SC/NC, with an incident alert generated and shared across state lines in real time.
5. **One federated corridor**
   Sovereign at the edge, federated at the center — that's OpenITS.

Source: Vikasa / OpenITS live corridor walkthrough (2026)

FOR YOUR SECURITY ARCHITECTS AND AUDIT LEADS

Open does not mean exposed. Trust is structural.

There is no shared cross-tier secret and no single authorization component to compromise. Every tier owns its own trust boundary; cross-tier reach is always a deliberate, subject-filtered subscription at a DMZ.

Mutual TLS everywhere

Verified identity at the TLS layer on every connection — no anonymous peers.

Tier-scoped identity

JWTs or NATS-native NKEYs, each scoped to a single tier. No credential is reusable across boundaries.

Per-connection subject ACLs

Authorization enforced on the terminating server. An agency shares exactly what it chooses — nothing more.

Seven isolated trust roles

Cabinet, regional DOT, DOT core, DOT DMZ, national/coalition DMZ, third-party, vendor — structurally separated.

vikasa

Source: Vikasa — trust & security architecture (docs/exec/03)

DESIGNED FOR DATA TO SCALE

Data arrives resolved making Digital Twins possible - at scale.

Our design uses edge compute at the source to write telemetry to files, MQTT, or TCP streams. Vikasa slots in cleanly: it normalizes multi-vendor edge data into shared schemas first, then feeds downstream systems through the very same ingestion paths — turning single-vendor point integrations into open, governed interoperability.

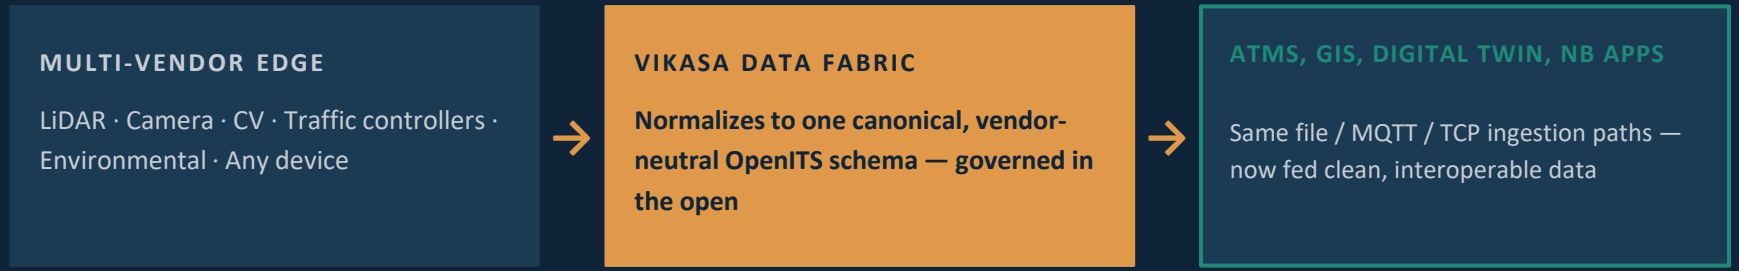

The result for your network — real-time processing, security, scale — stays intact. Vikasa adds the missing fourth pillar: interoperability, with digital twins reading the same canonical stream as every other consumer.

BUILT TO FEED WHAT YOU ALREADY RUN

Vikasa speaks the language every northbound platform already reads.

Your data needs to live and flow to any destination. Vikasa normalizes multi-vendor edge data into a single canonical, vendor-neutral event stream — then publishes it once to whichever traffic-ops, GIS, cloud-analytics, or asset-management platform your agency has already invested in. No rip-and-replace. No bespoke integration per application.

VIKASA CANONICAL, VENDOR-NEUTRAL DATA FABRIC

Multi-vendor edge data → one schema → published once

Traffic Ops (ATMS)

Lonestar™ ATMS
SunGuide
Kapsch EcoTrafiX

GIS & Asset Mgmt

Esri ArcGIS
Bentley AssetWise
ServiceNow

Cloud Data Platforms

Snowflake
Databricks
Power BI · Tableau

Digital Twins

Bentley iTwin
Esri ArcGIS Urban
Cesium

Security & SIEM

Splunk
Microsoft Sentinel
Other SIEM platforms

Same canonical stream, five different doors — Digital Twins read the identical feed as every other platform.

Source: TxDOT TSMO (Lonestar™ ATMS), Snowflake, Esri, Bentley Systems, Cisco Live BRKXAR-2031

WHERE THE TRANSPORTATION BUDGET ACTUALLY GOES

Time to connect versus time to innovate.

Today, agencies burn most of their technology budget just wiring systems together — every device, every dashboard, a new custom integration. Vikasa integrates the data once into an open canonical fabric, so that spend flips: the budget moves from plumbing to progress.

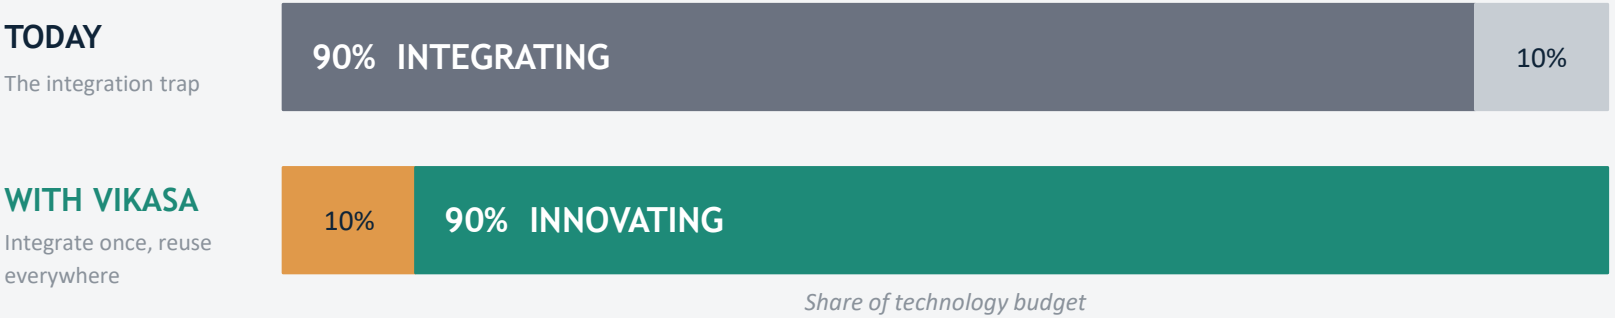

Integrate once

Normalize every vendor's data into one open schema — a single connection, not one per system.

Reuse everywhere

ATMS, GIS, cloud analytics, digital twins, and SIEM all read the same canonical stream.

Let the data move

Share across districts, agencies, and states without a new contract or a central broker.

Source: FHWA CodeFest Playbook (Deliverable D26), 2M Transformation Services & 360ns field experience

WHAT CANONICAL DATA MAKES POSSIBLE

Clean, shared data is the fuel every agency ambition runs on.

Unlock the power of AI at the edge: LSTM, GRU, and graph neural networks for congestion, incident detection, and trajectory prediction are only as good as the data feeding them. Interoperable, canonical data is what turns these from single-intersection demos into corridor- and state-wide capabilities — and what keeps a digital twin current instead of a static snapshot.

Predictive congestion management

AI trained on multi-agency, multi-vendor feeds — not one vendor's slice

Real-time incident detection

Shared incident, red-light-running, and overspeed topics across a corridor

Vulnerable road user protection

VRU detection events flowing to any consuming safety application

Semantically rich digital twins

Canonical real-time + lifecycle data feeding twin models that stay current, not static snapshots

Corridor & Vision Zero coordination

Multiple agencies acting on the same real-time picture toward zero fatalities

vikasa

THE ARGUMENT, IN ONE SLIDE

Why open-source. Why now.

Open is the only way to escape lock-in

A standard nobody controls is a standard every vendor can build on — and every agency can trust. Apache-2.0 stewardship, an operator-weighted steering committee, and a graduation path that lets vendors innovate without forking.

The federal groundwork is already laid

This work sits inside FHWA's Digital Infrastructure Enablers program and operationalizes its data-governance and use-case research. The research lineage is federal; the reference implementation is running.

The industry and the hardware are converging now

Tech is investing heavily in open-source models and industrial edge. Agencies and vendors are already designing next-generation ITS data platforms. The canonical layer that ties it together does not exist yet — and we have built it.

THE OPPORTUNITY IN FRONT OF THIS GROUP

A standard this group can help shape together.

The interoperability layer for American Transportation Data is being defined right now. This group can help shape it — with distributed edge compute, Vikasa's OpenITS fabric, and the platforms your agencies are already building.

1

Share what you're already building

Knowledge sharing / collaboration objective

2

Point to the lifecycle data gaps that matter most

Address needs objective

3

Help define the canonical schema

Encourage standardization objective

DDSG does not set policy or vote on positions — this is offered as a technical contribution for the group to build on.

vikasa

The infrastructure is public.

The data beneath it should be too.

Let's build the open data fabric for the transportation industry — together.

Vikasa · OpenITS · FHWA CodeFest · Stream Konnect

2M Transformation Services & 360ns