Report P5B Flood Data Services

Prepared for Project 0-7095-01 Flood Assessment System for TxDOT (FAST)

By Andy Carter, David Maidment and Tim Whiteaker
Center for Water and the Environment, University of Texas at Austin

Attila Bibok. Matt Ables
KISTERS North America, Inc.

31 July 2025

Introduction

Task 5 focuses on the development and deployment of Flood Data Services to support the functionality of FAST map applications. These services include (1) a geospatial representation of road segments affected by flooding, and (2) a representation of bridges with active flood warnings where water levels are approaching or exceeding the bridge's low chord elevation, which are each informed by (3), an estimate of the extent of flood inundation.

Roadflood.com

Roadflood.com is a web application developed by the University of Texas at Austin as part of Task 5, designed to demonstrate a prototype deployment of operational flood warning geospatial layers for TxDOT. Development began with a prototype deployment over the 55 square mile Mustang Creek watershed encompassing the City of Taylor, Texas. https://roadflood.com/fast_historic.html

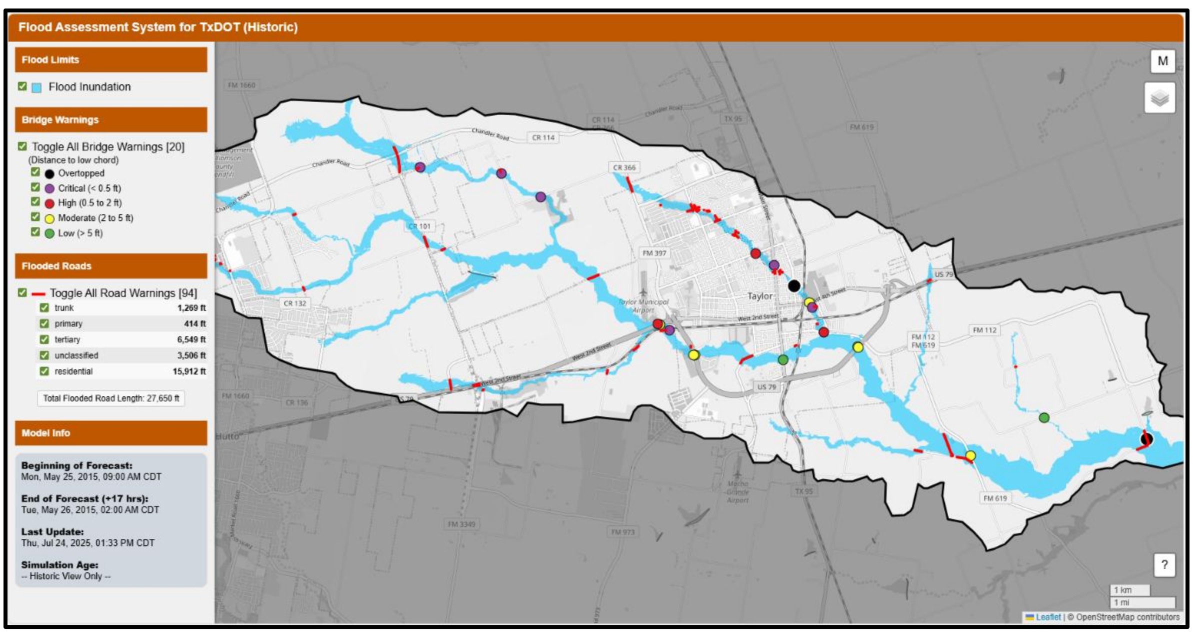

Figure 1. Taylor, TX – Roadflood.com prototype – mapping a historical flood (May 25, 2015)

On July 4, 2025, the Roadflood.com prototype was successfully scaled and deployed statewide, covering all of Texas. The system now provides hourly updates of flood extents across the state, based on hydrologic data from the National Water Model (NWM) short-range forecast. These flood limits reflect the maximum predicted flood levels within an 18-hour forecast window. Forecasts are generated independently for each of TxDOT's 25 districts, allowing for district-specific flood risk visualization and response planning. A status viewer is available at https://roadflood.com/fast_status_viewer.html , where users can see the latest update time for each district and access real-time flood warning maps.

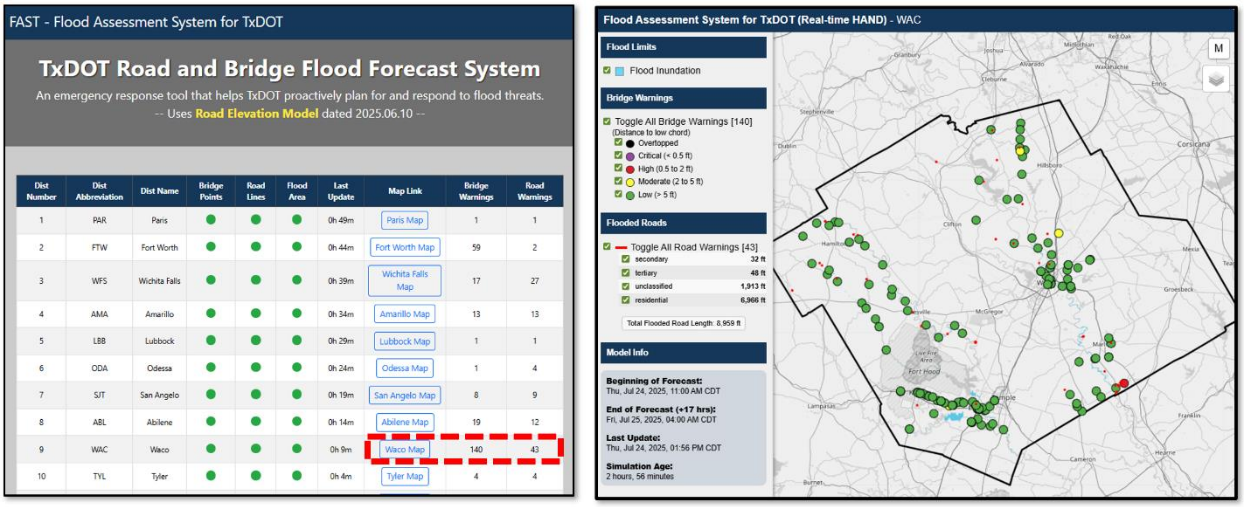

Figure 2. Roadflood.com - FAST status viewer and sample district map (Waco District)

The FAST flood warning system delivers data in three formats: (1) flood inundation polygons, (2) flooded roadway segments, and (3) bridge warning points. The Roadflood.com web interface was developed to demonstrate how these continuously updated datasets can be effectively visualized. Ultimately, TxDOT may choose to integrate these flood layers into platforms such as the Statewide Planning Map. Roadflood.com serves as a proof-of-concept, showcasing the capabilities of the statewide FAST flood warning system.

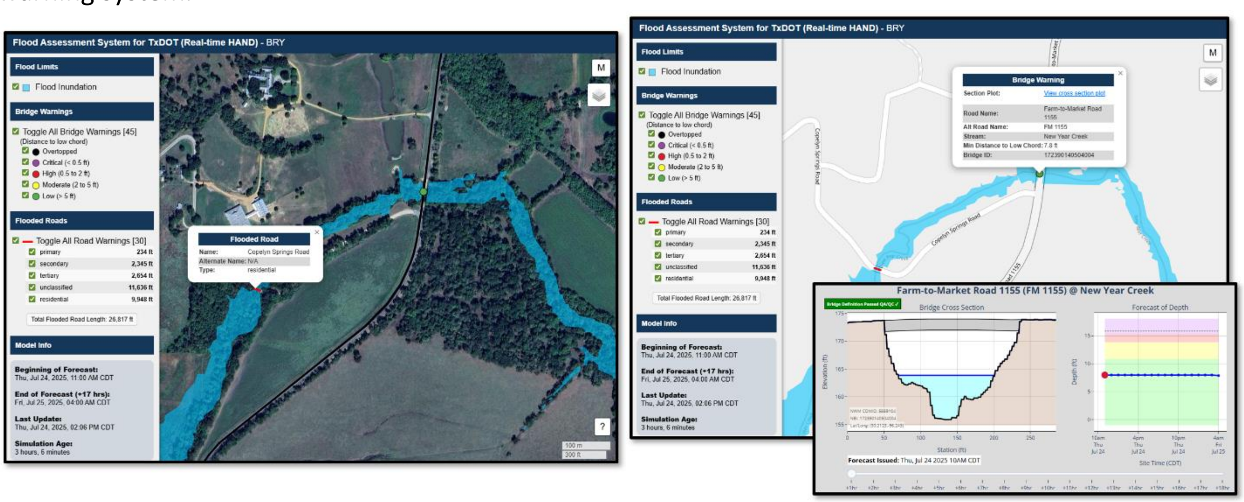

Figure 3. Roadflood.com interface samples – FM 1155 at New Years Creek

Road Elevation Model

The Road Elevation Model is a digital surface dataset representing roadway terrain, developed to support flood impact assessments along vehicular corridors. Created by Oak Ridge National Laboratory (ORNL) using LiDAR point cloud data, the model focuses on buffered zones around OpenStreetMap (OSM) road centerlines to capture the elevation of the actual travel surface. This dataset is essential for determining where forecasted floodwaters may intersect roads and for establishing both flood inundation polygons and flooded road discharge thresholds. Road elevation data for the entire state of Texas are available at: https://web.corral.tacc.utexas.edu/nfiedata/road3d/

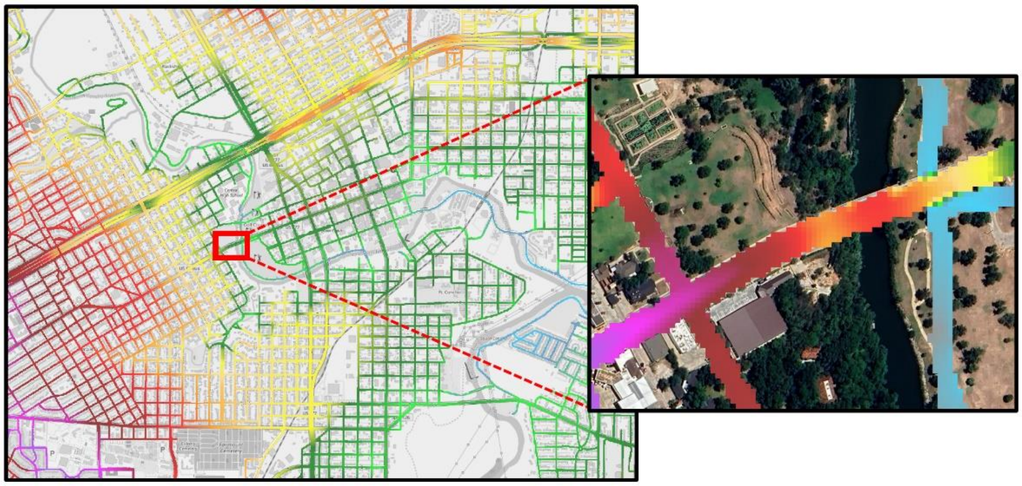

Figure 4. Road Elevation Model sample – City of San Angelo, TX – Inset West Beauregard Avenue @ North Concho River

Static Layers

Before the FAST system can operate in real-time, it requires a foundation of static geospatial and tabular data to support the generation of dynamic flood warnings. These static datasets are pre-computed and staged to establish the relationships between forecasted hydrologic conditions and their geospatial impacts. This pre-processing involves (1) generating flood inundation polygons based on terrain models, stream networks, and flood thresholds, (2) identifying roadway segments and bridge locations that intersect with predicted flood extents and (3) assigning thresholds and metadata (road type, functional class, bridge elevation) to inform impact assessments.

Flood Inundation Maps

Flood inundation polygons for various flow rates along each National Water Model stream in Texas were originally generated to support the Texas Department of Emergency Management (TDEM) as part of the "Pin2Flood" project. These polygons were created at approximately 1-foot vertical intervals using the Height Above Nearest Drainage (HAND) method applied to 3-meter terrain data. For integration into the FAST system, each polygon was further refined using the road elevation model to mask out road segments that remain unflooded (i.e., not overtopped) at a given flow rate.

Across the entire state, the FAST system includes more than 7.1 million precomputed flood polygons, each adjusted to reflect road elevation model terrain.

Figure 5. Sample Flood Inundation Polygon adjusted for Road Elevation Model - West Beauregard Avenue @ North Concho River

Flooded Roads

Using OpenStreetMap (OSM) road linework, the vehicular road network in Texas was divided into over 2.6 million individual segments. For each segment, a flood trigger flow value was identified representing the river flow at which that segment begins to show signs of flooding. This was determined by intersecting each road segment with precomputed flood inundation polygons, allowing the system to establish the flow threshold at which a specific road is overtopped.

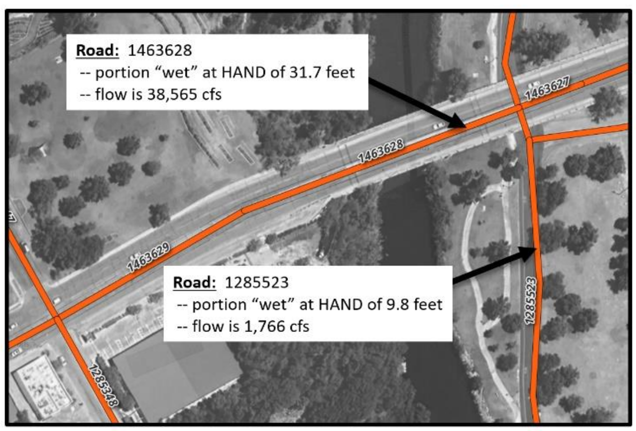

Figure 6. Sample Flood ed Road Thresholds - West Beauregard Avenue @ North Concho River

Bridge Warnings

Warning data for points for bridges that cross over a stream identified in the National Water Model (NWM) were extracted using the "TX-Bridge" methods detailed in section 5.2.1 of Technical Memorandum FY24RU1 [TMFY24RU1Project0_7095_01.pdf].

Streamflow Forecasts

FAST converts hydrologic forecasts into actionable flood mapping products by translating predicted stream flows into spatial representations of inundation, flooded roads, and bridge hazards. The FAST system relies on two main sources of hydrologic forecast data. They are (1) the National Water Model and (2) the data assimilation scheme developed by Dr. Matt Bartos of the University of Texas.

National Water Model

The National Water Model (NWM) output file, produced by NOAA's National Centers for Environmental Prediction (NCEP), provides forecast data in NetCDF format covering the most recent two days. Files are hosted at https://nomads.ncep.noaa.gov/pub/data/nccf/com/nwm/prod/ .

For example, a file named "`nwm.t03z.short_range.channel_rt.f011.conus.nc`" contains predicted streamflow values from the channel routing model for the continental United Stated (CONUS), generated for the 03Z (3 AM UTC) short-range forecast run, showing data for forecast hour 11 (i.e., 14:00 UTC). Each hourly short-range (i.e t03z) forecast run includes 18 separate forecast files (i.e., f001 to f018). NCEP stream flow forecast data are typically delayed by 1.5 to 2 hours behind the current time. This delay is from the time to run the hydrologic models and to publish the forecast outputs. As a result, the most recent forecast data available at any given moment usually corresponds to conditions from about two hours earlier.

The second source is the data assimilation model, which integrates real-time observations—such as stream gauge readings and precipitation data—into the modeling framework to produce a more accurate, observation-informed estimate of current hydrologic conditions. This data assimilation process improves near-term forecasts and provides a reliable baseline for initializing future NWM forecast cycles.

Data Assimilation

The goal of data assimilation (DA) is to improve the National Water Model (NWM) stream flow forecasts by using flow data from additional stream gages. These include the RQ30 radar gauges in addition to gages from regional authorities like the Lower Colorado River Authority (LCRA).

Because the NWM forecasts are delayed by about two hours, the earliest forecast data is already outdated by the time it's received. To improve this, a data assimilation (DA) algorithm is applied to update the forecasts using the latest observations.

When a new short-range forecast is released, streamflow measurements from various sources are gathered through a system known as Datasphere which is hosted by KISTERS. These observations are then used to adjust the NWM raw forecast, producing a more accurate version of current conditions. While the original NWM forecast includes observations that are nearly 2 hours old, the DA forecast can provide a fresher forecast only 45 minutes old, depending on the location of available gages.

Adjusted stream flows from DA runs are automatically computed within KISTERS Datasphere. Operational files are located at https://knatempstorage.s3.us-west-1.amazonaws.com/nwm_txdot_output/Readme.html .

Time Series Processing

FAST uses streamflow data from various gage types (including RQ-30 radar sensors) to improve forecast accuracy. Since these gages differ in how and when they report data, their readings must be harmonized before being used in the data assimilation algorithm. KISTERS collects and integrates these readings into their Datasphere system, which is designed to handle multiple data sources. This setup allows FAST to process stream flow observations in near real-time.

The 80 “Streamflow 1” RQ-30 gages report their data with a delay of typically 2 to 5 minutes because they’re directly connected to KISTERS. However, most other gages transmit less frequently in sometimes in batched bursts. Datasphere tracks how delayed each gage’s data is and can flag any that fall behind expected timing.

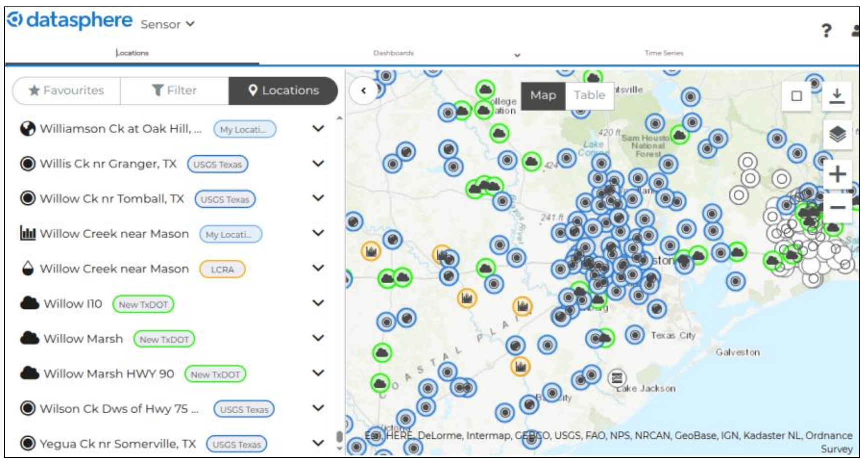

Figure 7. Datasphere – Example showing multiple gages around Houston, TX

Backend Service

The FAST system transforms real-time hydrologic forecast data into actionable flood mapping products through an automated, cloud-based process. This system is built on Amazon Web Services (AWS) infrastructure, with several coordinated components that work together to ingest forecast data, process it against known flood thresholds, and publish updated map layers for use by downstream applications.

At the heart of the system is a Python application deployed on an Amazon EC2 instance. This EC2 machine runs a Docker container that holds the FAST application code, which is publicly available at github.com/andycarter-pe/fast_realtime. The FAST application monitors new short-range hydrologic forecasts and, when detected, triggers a SQL script. This script performs spatial queries and calculations by connecting to a PostgreSQL/PostGIS database hosted on Amazon RDS.

The database holds static flood-related layers for each stream segment. These include flood inundation polygons, road flooding thresholds, bridge overtopping points, and rating curves. When new streamflow forecasts arrive, the SQL query compares the forecasted flows to these thresholds to determine whether flooding is likely. This process results in the generation of several “dynamic” layers: flooded road lines, active bridge warning points, and flood inundation polygons that reflect the latest hydrologic forecast.

Once generated, these spatial layers are exported from the database in GeoJSON format. The resulting files are uploaded to an Amazon S3 bucket at s3://txdot-realtime-hand-v2, where they are organized by TxDOT district for easy access. These files are intended for use in GIS software platforms like ArcGIS or QGIS, and are also ingested by custom web applications. One such example is Roadflood.com, a public-facing web platform that reads the GeoJSON outputs and presents an interactive map showing flood conditions across Texas. A district-specific example of this viewer is available here for Waco District, where users can explore real-time warnings for roads and bridges.

By automating this end-to-end workflow, the FAST system provides near real-time flood risk information with minimal manual intervention. It allows TxDOT to access location-specific flood impact data as hydrologic conditions evolve.

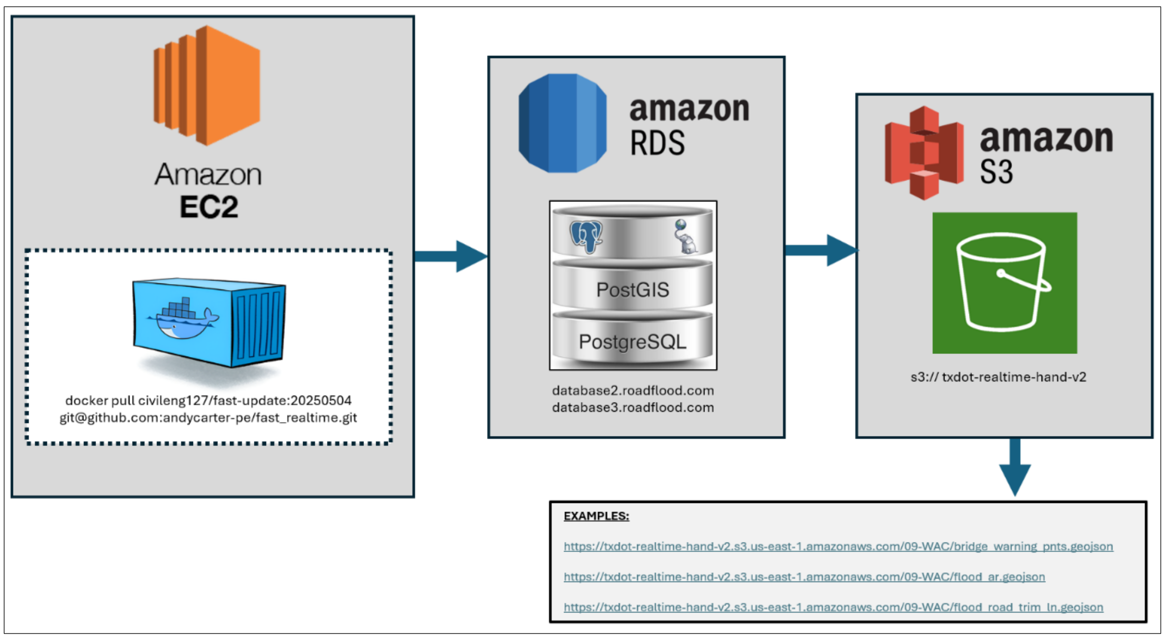

Figure 8. FAST Backend Service Diagram

The FAST system's data infrastructure is built using PostgreSQL, an open-source relational database, enhanced with PostGIS to support spatial queries and geographic data types. To manage flood mapping at a regional level, a separate FAST database instance was developed for each TxDOT district. PostgreSQL database dumps for each TxDOT district are publicly available at: https://web.corral.tacc.utexas.edu/nfiedata/acarter/roadflood-db-20250702/. Combined, the full set of FAST databases for all TxDOT districts represents approximately 63 GB of data.

KISTERS Deployment

The FAST system, developed and deployed on Amazon Web Services (AWS), has been maintained by a research team at the University of Texas as of July 2025. While this academic deployment has been effective for testing operations, the goal is to transition the system to a professionally managed cloud environment operated by KISTERS by the end of August 2025. KISTERS brings enterprise-level expertise in cloud infrastructure, including real-time error monitoring, automated scaling, and guaranteed up time.

In addition to technical robustness, KISTERS' infrastructure provides support for data assimilation workflows and the integration of real-time streamflow observations. This enables more accurate and responsive flood forecasting, especially in areas vulnerable to flash flooding. Their platform also supports secure access, routine backups, and mirrored databases for redundancy and disaster recovery.

Migrating FAST to KISTERS marks a shift from a research prototype to a production-grade, enterprise-ready system that can reliably support statewide operations.