



Shreveport Section

Virtual Talk

Tuesday, October 13 at 6:00 PM

Free & Open to All

Registration Required

AI Observability for Reliable Production Systems

An SRE Perspective



Ravi Teja Reddy Mandala

*Senior Site Reliability Engineer
Oracle America Inc.*



ABSTRACT

As artificial intelligence becomes embedded in production operations, conventional observability practices are no longer sufficient to explain system behavior, distinguish meaningful anomalies from noise, or support timely and accountable decisions. This talk presents AI observability from a site reliability engineering perspective, focusing on how telemetry, machine learning, and operational context can be combined to improve the reliability of large-scale production systems. It examines the signals required to observe both AI-enabled infrastructure and the AI components used within operational workflows, including metrics, logs, traces, model outputs, data quality, drift, confidence, and human interventions.

The session introduces a practical lifecycle for converting these signals into dependable SRE outcomes: establishing service objectives, detecting abnormal behavior, correlating symptoms across distributed components, prioritizing incidents, supporting root-cause analysis, and validating recovery. It also addresses common production risks such as false positives, opaque model decisions, automation bias, incomplete telemetry, privacy concerns, and unbounded remediation. Particular attention is given to the safeguards needed before AI can influence incident response, including explainability, confidence thresholds, human approval, audit trails, rollback mechanisms, and continuous evaluation. Attendees will leave with a structured approach for introducing AI into observability without weakening operational discipline, and with practical guidance for designing systems that are measurable, resilient, and trustworthy in production.