Frequently Asked Questions

This document provides additional context for the TRACE Framework, which is a practical guide for monitoring AI-enabled clinical decision support (AI-CDS) tools after they have been deployed in clinical settings. It is intended for developers, implementers, and governance bodies considering or preparing to adopt TRACE.

What is TRACE?

TRACE is a framework for postdeployment monitoring of AI-enabled clinical decision support tools. It provides practical guidance for identifying risks, monitoring how tools are used, and ensuring that AI-CDS systems remain safe, effective, and equitable over time. TRACE covers 5 key areas: Technical Integrity, Real-World Use, Alignment and Accuracy, Clinical Fairness, and Escalation and Safety Response. It is intended to help AI developers and implementers track the performance of AI-CDS tools in real-world settings.

How was TRACE created?

TRACE was developed based on field experience deploying AI-CDS tools, a review of existing research on AI safety and governance, and interviews with stakeholders across healthcare systems, industry, and oversight bodies. The framework reflects real-world challenges and needs, and was designed to translate high-level principles into practical monitoring strategies that can be applied across diverse clinical settings.

Who should use TRACE?

Developers building AI-CDS tools for healthcare environments;
Health systems deploying and managing AI-CDS tools; and
Governance bodies such as clinical AI oversight committees or patient safety boards.

Does TRACE address privacy?

TRACE is not a privacy compliance framework. However, many of the data elements used for monitoring – such as patient demographics, clinician identifiers, and encounter details – are protected health information. All data collection, storage, and analysis conducted as part of TRACE implementation must comply with applicable privacy regulations (e.g., the Health Insurance Portability and Accountability Act in the United States or the General Data Protection Regulation in the European Union) and institutional data governance policies. Privacy and security safeguards should be embedded in the governance structures that oversee TRACE implementation.

ASPEN POLICY ACADEMY
aspen institute

How does TRACE fit with existing regulations and standards?

TRACE is designed to complement existing regulatory and safety frameworks. It can be aligned with the following:

- Postmarket surveillance requirements of the US Food and Drug Administration (where applicable);
- The Blueprint for Trustworthy AI Implementation Guidance and Assurance for Healthcare of the Coalition for Health AI (CHAI);
- The AI Risk Management Framework of the National Institute of Standards and Technology (NIST); and
- Organizations' internal quality and safety programs.

Using TRACE helps organizations operationalize the guidelines from bodies such as CHAI and NIST, providing a practical approach to monitoring.

What is AI-CDS?	What is post-deployment monitoring?	Does TRACE require building custom instrumentation and monitoring tools?
AI-enabled clinical decision support refers to tools that use artificial intelligence to aid clinicians in making decisions about patient diagnosis, treatment, triage, or care planning. These tools are designed to support clinical judgment and are typically integrated into existing workflows such as the electronic health record or point-of-care systems.	Post-deployment monitoring is the ongoing evaluation of a tool's performance, safety, and equity once it is in use. It addresses the reality that patient populations, workflows, and AI model behavior can change over time, sometimes reducing effectiveness or introducing risk.	Not necessarily. TRACE is technology-agnostic and can be implemented using either in-house systems or third-party vendor solutions. The framework defines what to monitor and how, but not which tools to use. Organizations may choose to build custom telemetry and analytic pipelines, adapt existing quality improvement systems, or contract with vendors that provide monitoring capabilities.

ASPEN POLICY ACADEMY
aspen institute

How often should monitoring occur?	Can TRACE be applied to tools that are already deployed?	Is TRACE specific to large language models?
Monitoring frequency depends on the tool's risk profile, scope of deployment, and rate of change. High-risk tools or those in dynamic clinical domains may require continuous or weekly monitoring, while lower-risk tools may be reviewed monthly or quarterly.	Yes. TRACE can be implemented retrospectively, provided that the necessary data sources are available. In some cases, additional logging or telemetry may need to be added to support full monitoring.	No. TRACE is designed to support monitoring for all types of AI-powered clinical decision support tools, not just those using large language models (LLMs). While it includes guidance for risks that are more common with generative AI (including LLMs), TRACE can also be used for tools built on traditional machine learning approaches, rules-based systems, or combinations of different methods.

What are some potential use cases for applying TRACE?

TRACE can be applied to AI-CDS tools that assist clinicians in diagnosis, triage, and treatment planning. The following are some hypothetical examples:

1. A hospital regularly reviews the outputs of an AI tool that suggests possible diagnoses to ensure that they remain consistent with current clinical guidelines and have not degraded over time.
2. An urgent care clinic monitors how an AI triage tool performs across different patient groups, checking for disparities in recommended care pathways based on race, language, or insurance status.
3. A health system detects that clinicians are overriding an AI tool that recommends antibiotics more often than expected. This prompts a review to see if the tool's recommendations are outdated or out of line with the way clinicians are actually treating patients.
4. A cancer center flags that a treatment recommendation tool is no longer aligned with updated oncology guidelines, leading to a review and update process.

If you'd like to learn more, see the full project, including a policy brief explaining the fellow's core recommendations, at aspenpolicyacademy.org/project/ai-cds-monitoring-2025.