



ONC 45 CFR 170.315(B)(11)(VI)

Summary of Intervention Risk Management Practices

Predictive Decision Support Intervention: AI Assistant Chatbot (Ambulatory)

Prepared by: Azalea Health | Regulatory Compliance

Date: August 18, 2026

Field	Value
Classification	PUBLIC
Status	FINAL
Version	1.0
Date	2026-08-18
Regulatory basis	45 CFR 170.315(b)(11)(vi) — Intervention risk management
Certification ID (Azalea EHR — ambulatory)	15.04.04.2688.Azal.04.02.1.221227

1. Overview

1.1 Purpose

This document outlines the steps and controls Azalea Health has implemented to ensure the responsible development, deployment, and use of certain AI-powered features within its Electronic Health Record (EHR) system and associated components. It summarizes risk management activities performed by Azalea Health for its Predictive Decision Support Intervention(s) (Predictive DSI) in compliance with the ONC Health IT Certification Program.

A Predictive DSI is technology that supports decision-making based on algorithms or models that derive relationships from training data and then produce an output resulting in prediction, classification, recommendation, evaluation, or analysis. These features help leverage data and predictive analytics to aid healthcare decision-making.

Azalea Health's commitment to safety, privacy, transparency, and preventing bias remains central as AI technology is integrated into its products.

1.2 Scope

This document covers the AI Assistant Chatbot, a conversational assistant available to clinicians and staff within the Azalea ambulatory EHR (certification ID 15.04.04.2688.Azal.04.02.1.221227). It does not cover any other Azalea product or AI feature.

2. Terms

Term	Definition
Predictive Decision Support Intervention (Predictive DSI)	Per ASTP/ONC: technology that supports decision-making based on algorithms or models that derive relationships from training data and then produce an output that results in prediction, classification, recommendation, evaluation, or analysis. Not all Azalea AI features meet this definition.
AI Assistant Chatbot	Azalea's conversational assistant for ambulatory clinicians and staff, providing chart lookup, clinical-note retrieval, knowledge-article search, and coding support within a chat interface.

3. Intervention Risk Management Plan

The following risk management activities are performed as part of Azalea Health's risk management processes.

3.1 Intervention Description

The AI Assistant Chatbot is a conversational, read-only assistant for ambulatory clinicians and staff. It is always invoked in the context of a single, specific patient; there is no path by which a user can launch the assistant outside of a patient's record. Within that patient-scoped session it answers questions about the patient's medical history, medications, allergies, diagnoses, immunizations, and the current encounter's clinical notes and findings, and supports natural follow-up questions within a conversation. The assistant can additionally answer general Azalea product and workflow questions grounded in Azalea's published customer help content; this is a capability available within the patient-scoped session, not a separate mode of operation.

All output is advisory and draft-only: the assistant does not write to the chart, place orders, change settings, or take any other automated action. Every response is intended to be reviewed by the user before being relied upon. The assistant is not a diagnostic tool. It does not generate patient-specific diagnoses, treatment recommendations, or clinical decision support, and does not replace clinical judgment.

3.2 Process Overview

When Azalea Health determines that an AI feature meets the definition of a Predictive DSI, the intervention is documented and analyzed to identify risks, including those related to the validity, reliability, robustness, fairness, intelligibility, safety, security, and privacy of its output. Each Predictive DSI is subject to risk analysis and mitigation activities commensurate with its risk profile, as outlined below.

3.2.1 Risk Analysis

Intelligibility. The AI Assistant Chatbot displays a persistent notice, visible for the duration of the chat session, informing the user that responses are AI-generated, that chat sessions are private to the user, and that submitted data is not used to train Azalea's AI models. Users are directed to review AI-generated output before relying on it in patient care or billing-related decisions.

Security. The AI Assistant Chatbot is subject to Azalea Health's software security program, including secure code review, automated testing, dependency and vulnerability scanning, and a defined control set for cloud-hosted AI models, including account and region isolation, least-privilege access, encryption, logging, and monitoring. Azalea Health maintains an incident-escalation process for security concerns identified in any AI feature, including this one.

Reliability and Robustness. Prior to broader release, the AI Assistant Chatbot is made available to a limited set of beta clients operating in their production environment, providing a practical setting in which to evaluate the intervention's performance against real patient data and clinical workflows. This testing surfaces edge cases and operational issues not captured during development and informs ongoing reliability improvements. Azalea Health monitors production usage on an ongoing basis, including response completion and error rates, to evaluate reliable operation over time.

Fairness. Inputs to the AI Assistant Chatbot are limited to the clinical, coding, and knowledge-article information necessary for the intervention to function, retrieved through scoped queries rather than an open demographic data feed. Azalea Health maintains an ongoing fairness review for this intervention. Because the intervention retrieves only the record of the patient the user has already opened, Azalea evaluates fairness at the level of intervention behavior rather than model inputs: production usage is monitored in aggregate to compare how consistently the intervention responds across client sites, specialties, and user roles, and to identify any group for which responses are disproportionately incomplete, unanswered, or in error. That review is conducted on a recurring basis by Azalea's product and clinical stakeholders, and any disparity it surfaces is documented and routed through the risk mitigation and governance processes described in Sections 3.2.2 and 3.2.3. Fairness findings are also reassessed whenever the intervention's model, data sources, prompts, or intended use materially change.

Privacy. The AI Assistant Chatbot operates within Azalea Health's HIPAA-compliant infrastructure. The underlying model provider is contractually prohibited from using submitted data to train or improve its models, and inputs and outputs are not retained by the model provider. Use of the assistant is governed by Azalea Health's Business Associate Agreements, data classification policies, and minimum-necessary-use principles, and chat data is subject to defined retention limits.

Validity and Safety. The AI Assistant Chatbot is reviewed by Azalea's Product Management and Quality Assurance teams as part of standard software validation. Its output is advisory only. The assistant does not write to the patient chart, place orders, or take automated clinical or billing action, and its output is subject to review by the treating clinician or staff member before being relied upon. Azalea Health maintains an intervention-specific hazard analysis for the AI Assistant Chatbot, consistent with recognized clinical risk management practices. That analysis identifies the intervention's foreseeable failure modes, including incomplete, inaccurate, outdated, or out-of-context responses. It rates each for clinical severity and likelihood, and maps each to the specific control that mitigates it, such as the advisory-only output model, the requirement for clinician review, patient-scoped retrieval, and the persistent AI-generated-content notice. The hazard analysis and the associated clinical-safety review are conducted by Azalea's Product Management and

Quality Assurance teams with clinical input, are a prerequisite to release, and are revisited when the intervention's design, model, data sources, or intended use materially change, or when post-release monitoring or user-reported issues indicate a new or changed risk.

3.2.2 Risk Mitigation

Where risk analysis identifies a risk associated with the AI Assistant Chatbot, Azalea Health applies mitigating controls through product design (for example, limiting the intervention to read-only, advisory functionality) or through informational controls that direct the user to take an appropriate action, such as the AI-generated content notice described above. Azalea Health's AI governance policies define mandatory security, privacy, and human-review requirements applicable to this intervention, and Azalea's Terms of Use state that accepting AI-generated output without review is a violation of those Terms.

3.2.3 Governance

All activities described in this risk management plan, and the design, development, and maintenance of the AI Assistant Chatbot generally, are subject to Azalea Health's AI governance policy, information security management program, and software development lifecycle standards. Azalea Health maintains an inventory of approved AI tools and models in active use, reviewed on a defined periodic basis, and assigns AI governance responsibilities across Security, Compliance, IT, Legal, Product & Engineering, and business owners. Azalea Health's Terms of Use additionally commit to maintaining a list of the AI model providers in use and to logging use of AI Features.