Inteca | IT's about business

Security Overview

PractIQ - AI Delivery Operating System

Inteca | Red Hat

Introduction

The adoption of AI in enterprise environments introduces new security requirements. AI agents operate on organizational data, invoke external language models, and take autonomous actions on behalf of users—introducing new risk classes that traditional access controls were never designed to address.

PractIQ is built for enterprise environments with stringent security requirements—particularly in the financial, banking, and insurance sectors. Security is not an add-on in PractIQ; it is embedded in the platform's architecture at every level: from runtime isolation, to data access control, to protection against uncontrolled information flows to language models.

This document describes the key security principles and mechanisms underpinning the PractIQ platform.

Key security principles

1. Runtime environment isolation

Each practice running in PractIQ executes within a dedicated, tightly controlled container environment on Kubernetes—in the recommended production configuration on Red Hat OpenShift. The environment for each practice is predefined and immutable during execution: agents have access only to the tools, data, and services required to carry out that specific practice. No resources outside the defined scope are accessible within the execution environment.

Environment isolation eliminates the risk of unauthorized cross-practice access and ensures full predictability of agent behavior—regardless of how many users and agents are running concurrently on the platform.

Inteca Sp. z o.o.
Powstańców Śląskich 9
53-332 Wrocław
+48 71 715 60 91
biuro@inteca.pl
inteca.com
practiq.tech

2. Zero Trust and Attribute-Based Access Control

Data access—for both users and AI agents—is governed by a Zero Trust model with Attribute-Based Access Control (ABAC). This means that permissions are not granted based on the user's role alone, but on a combination of attributes: identity, the context of the practice being executed, data classification, and current session state.

In practice: an AI agent executing a requirements analysis practice has access only to the data and tools defined for that practice—it cannot access data associated with another practice or invoke services outside its defined privilege scope. This approach is access-method-agnostic: it applies equally to Model Context Protocol (MCP) calls and direct service invocations.

3. LLM gateway and guardrails

All language model calls pass through a central LLM gateway built on LiteLLM, equipped with a set of control mechanisms (pre-call hooks) that are triggered before any data is sent to the model. These mechanisms scan and validate input, ensuring that no data failing to meet the defined security criteria reaches an external or local language model.

The LLM gateway serves as a central point of control and audit for all agent interactions with language models—enabling end-to-end logging, monitoring, and anomaly detection.

4. Local data classification and Data Loss Prevention (DLP)

PractIQ implements a local Data Loss Prevention (DLP) mechanism based on a domain ontology—not text patterns. Each practice operates exclusively on data defined in the organizational ontology. Data in the knowledge graph is classified by sensitivity level, and the local DLP gateway applies this classification to automatically mask sensitive data before it is transmitted to the language model.

The ontology-based approach—unlike classic regex-based DLP systems—delivers consistent, context-aware data masking across calls, eliminates false positives inherent to pattern matching, and operates entirely on-premises, without adding load to the LLM gateway and without sending data to external classification services.

Inteca Sp. z o.o.
Powstańców Śląskich 9
53-332 Wrocław
+48 71 715 60 91
biuro@inteca.pl
inteca.com
practiq.tech

Security in the context of Red Hat OpenShift

The recommended production environment for PractIQ is Red Hat OpenShift, which provides additional security layers at the infrastructure level. Built-in Role-Based Access Control (RBAC) and network policies enforce network isolation between platform components. Automated container image scanning detects known vulnerabilities prior to deployment. Cluster-level audit logging delivers a complete audit trail for all operations. Integration with Red Hat Build of Keycloak enables enterprise-grade identity and access management (SSO, MFA, identity federation).

Deploying PractIQ on-premises within an OpenShift cluster—recommended for organizations with the most stringent security requirements, such as banks and financial institutions—ensures full data sovereignty: no organizational data leaves the customer's infrastructure.

Summary

PractIQ's security approach is built on four mutually reinforcing principles: runtime environment isolation, attribute-based Zero Trust access control, a central LLM gateway with guardrails, and ontology-based local data classification and loss prevention. Combined with Red Hat OpenShift's security capabilities, these principles form a multi-layered security architecture that ensures AI agents operate strictly within defined boundaries—with no possibility of unauthorized data access, uncontrolled calls to language models, or leakage of sensitive information.

More info: inteca.com | practiq.tech

Inteca Sp. z o.o.
Powstańców Śląskich 9
53-332 Wrocław
+48 71 715 60 91
biuro@inteca.pl
inteca.com
practiq.tech