Inteca | IT's about business

Governed AI Delivery with continuous validation

The Case for an AI Delivery Operating System

How enterprise IT departments can overcome the governance gap and the legacy of the Agile era — and build a repeatable, scalable AI delivery model with PractIQ and Red Hat OpenShift.

Inteca | Red Hat

"AI does not scale by adding tools. It scales by building the operating model that governs how humans and AI agents work together — across new projects, existing systems, and legacy modernization."

Inteca Sp. z o.o.
Powstańców Śląskich 9
53-332 Wrocław

+48 71 715 60 91
biuro@inteca.pl

inteca.com
practiq.tech

Table of contents:

Inteca Sp. z o.o.
Powstańców Śląskich 9
53-332 Wrocław

+48 71 715 60 91
biuro@inteca.pl

inteca.com
practiq.tech

Executive summary

IT departments of large enterprise organizations are investing in artificial intelligence — but most of them are not able to translate these investments into systemic, predictable results. AI tools increase the productivity of individual engineers, but they don't solve a deeper problem: the lack of a unified model that enables AI people and agents to collaborate effectively across the organization.

However, the problem has two interrelated dimensions. The first is the lack of governance and standards for the new reality: organizations don't know how to implement AI in a controlled, repeatable, and scalable way — how to define agent work rules, standardize workflows, and ensure consistency of results across projects and teams.

The second is the legacy of two decades of work in agile methodologies. The Agile era has brought speed and iteration—but at the expense of architecture, documentation, and standardization. Organizations have built hundreds of applications without consistent technical documentation, with distributed and uncodified architectural knowledge, and with accumulating technological debt. As AI enters the manufacturing process, this debt becomes a critical barrier: AI agents can't work effectively without context — without knowing what the system does, how it's built, and what dependencies it has. Organizations that want to implement AI systemically must first rebuild what has been overlooked for years.

Without clear standards, rules, structured knowledge, documentation of AI-ready processes — AI remains an experiment rather than a strategic tool for a competitive advantage.

This document describes a new category of solutions—the AI Delivery Operating System—and shows how Inteca's PractIQ puts this concept into practice by running on Red Hat OpenShift. It is addressed to enterprise architects, technical directors and IT leaders responsible for the strategy of AI adoption in large organizations.

The key conclusions are as follows. First, the main barrier to systemic adoption of AI in IT departments is not the lack of technology, but the lack of governance and

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

an operational model of working with AI agents. Second, existing market solutions address only fragments of this problem – there is a lack of a platform that integrates governance, standardization of SDLC processes, agent orchestration, and automation of IT practices in one place. PractIQ with Red Hat OpenShift fills these gaps in three scenarios: for new projects (Greenfield), for existing systems that require documentation and architecture rebuild (Brownfield), and for legacy modernization and migration (Modernization) — providing a complete environment for systemic AI adoption regardless of your organization's starting point.

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

PART I: MARKET CONTEXT AND NEW CATEGORY

How software development is changing

Software development has undergone three fundamental transformations over the past three decades. Each of them required not only new tools, but above all a new operating model — a new way of thinking about how people, processes, and technologies work together.

The first transformation is the transition from waterfall models to agile methodologies. Waterfall assumed that the requirements could be defined in their entirety at the beginning of the project and then implemented sequentially over months. Agile challenged this paradigm — introducing iterativeness, short sprints, close customer engagement, and a willingness to change. The new operating model required new roles, new ceremonies, and a new way of thinking about planning. Organizations that have adopted it have gained speed and flexibility. Organizations that ignored it gradually lost their ability to compete.

The second transformation is the extension of the Agile paradigm to the area of implementations - DevOps. Agile methodologies have solved the problem of planning and iteration in software development, but they have left unsolved the problem of configuration management and deployments—how to efficiently and securely deliver software to production environments.? DevOps has answered this question by introducing CI/CD automation, infrastructure as code, a culture of collaboration between development and operations, and continuous monitoring. Delivery cycles have shortened from weeks to days and hours. However, one element remained unchanged: it was people who wrote the code, created the documentation, designed the architecture, and performed the tests.

The third transformation — the one we are witnessing today — changes this very element. Artificial intelligence ceases to be a tool supporting the work of engineers and becomes an active participant in the production process. AI agents analyze requirements, create technical documentation, design architecture, generate code, and perform tests. According to Gartner research, by 2028, 75% of software engineers will be using AI tools for coding, up from less than

10% in 2023. McKinsey estimates that generative AI can automate up to 30% of activity in typical software development projects.

However, this data only describes the scale of the phenomenon — it does not answer the question that is crucial for IT leaders today: how do we make this change bring systemic benefits to the organization, and not just a local improvement in the productivity of selected engineers?

History gives a clear answer: each of the previous transformations required a new operating model. Agile required new roles, ceremonies, and approaches to planning. DevOps required new tools, a culture of collaboration, and pipeline automation. The AI era is no exception – it requires a new approach to governance, standardization, and managing human and agent collaboration. Organizations that build this model first will gain an advantage analogous to that of the companies that were the first to implement DevOps.

Why AI adoption is stalling

IT departments of large enterprise-class organizations are facing a real challenge today: how to implement AI in a systemic way, solve governance problems that have accumulated over the years, and make AI work repetitively and predictably across the organization?

Contrary to appearances, the barrier is not technological. AI tools are available, language models are achieving impressive results, and the cost of using them is steadily decreasing. The problem lies elsewhere. Observations from Inteca projects in enterprise environments indicate five interrelated patterns that block systemic adoption.

Fragmented adoption and lack of standards. Different teams adopt AI independently of each other, using different tools, different models, and different approaches. Isolated islands of efficiency are being created – but the lack of common standards makes it impossible to replicate success at the scale of the organization. IDC reports that 64% of organizations cite the lack of consistent standards and processes as a major barrier to scaling AI initiatives.

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

Architecture and data not prepared for AI. AI agents need context—up-to-date technical documentation, consistent system architecture, structured domain knowledge. Meanwhile, in most large organizations, documentation is outdated or fragmented, architecture is inconsistent, and knowledge is dispersed between people and systems. Without an AI-ready knowledge base, agents operate in a vacuum—generating code that doesn't conform to the architecture, repeating errors, and requiring constant revision by engineers.

Agentic coding without governance. Agentic coding—a model in which AI agents perform sequences of programming tasks on their own—increases the speed of production, but at the same time introduces a new kind of unpredictability. Without clearly defined rules, coding standards, and validation mechanisms, every project looks different. The quality of AI-generated code is variable, costs rise unexpectedly, and the risk of technological debt increases instead of decreases.

Lack of scalability. What works in a single pilot team rarely replicates across the entire IT department. Successes remain local, because there is no platform that would allow to standardize and distribute proven practices of working with AI among projects, teams and locations. According to a 2024 Forrester report, only 17% of organizations that implemented AI pilots in IT were able to successfully scale them beyond the original team.

Debt inherited from the Agile era. The Agile era has brought speed and iteration—but at the expense of architecture, documentation, and standardization. For two decades, organizations have built hundreds of applications without consistent technical documentation, with architectural knowledge dispersed and uncodified, and with accumulating technological debt. When AI enters the manufacturing process, this debt becomes a critical barrier. As Craig Adam puts it in *"Agile is Out, Architecture is Back"* (Medium, 2025): *"Agile taught us speed. AI requires architecture."* AI agents can't work effectively without context — without knowing what the system does, how it's built, and what dependencies it has. Organizations that want to implement AI must first rebuild what they lost in the Agile era.

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

The business consequences of these problems are measurable. Longer time-to-market despite growing investments in AI. Inconsistent product quality and rising revision costs. Inability to demonstrate ROI from AI initiatives. And, perhaps the most serious long-term consequence, the increasing risk of losing control of the manufacturing process as AI becomes a key player.

Why existing solutions are not enough

The market for AI tools for IT departments is developed and growing dynamically today. However, the analysis of the available solutions reveals a structural gap: each of them addresses a selected part of the problem, none of them offers a systemic approach to governance and automation of the entire SDLC.

AI tools for developers—such as GitHub Copilot, Cursor, Claude Code, and Codex—effectively increase the productivity of individual developers when writing code. However, they do not offer any standardization, governance, or collaboration management mechanism at the scale of a team or organization. They are tools for individuals, not for organizations.

Enterprise-class automation platforms – UiPath, SS&C Blue Prism – focus on automating business and administrative processes. They are not designed with SDLC in mind, do not understand the specifics of the work of engineers, and do not offer governance mechanisms for AI agents working on code, architecture, or technical documentation.

AI governance platforms—IBM Watsonx as a leading example—offer powerful tools for AI model lifecycle management, monitoring, and compliance. However, they do not touch the process layer – they do not standardize the way AI works, they do not automate SDLC and they do not provide orchestration of agents in the production process.

Agentic AI platforms – such as Langchain – focus on agent orchestration and task automation. However, they lack a layer of standardization of IT practices, preparation of documentation and architecture for AI, and an end-to-end approach to the entire SDLC.

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

The result of this fragmentation is predictable: organizations that try to assemble a complete solution from the available point tools face problems with integration, consistency, and governance. No one manages the whole. The bonding layer is missing – a solution that organizes the way IT works in the AI era: it standardizes practices, manages agents, automates SDLC, and ensures governance at every stage of the manufacturing process.

New category: AI Delivery Operating System

Just as DevOps has defined a new operating model in the era of continuous integration and continuous delivery, the era of AI requires its own platform category. We propose to call it AI Delivery Operating System — AI-DOS for short.

The AI Delivery Operating System is an operational platform that standardizes, manages, and automates the way AI IT departments work throughout the software lifecycle. The analogy with the operating system is intentional here: just as the OS manages a computer's resources and enables applications to run in a predictable, controlled way, AI-DOS manages the collaboration of people, AI agents, data, and processes — ensuring that the whole thing works consistently, securely, and scalably.

The AI Delivery Operating System integrates four layers that have so far functioned independently of each other. The governance layer defines the standards and rules for working with AI agents – what they can do, how they should work, and how they are validated. The process layer provides structured SDLC workflows based on industry best practices—from requirements analysis to acceptance testing. The agent orchestration layer manages human and AI collaboration across workflows—assigning roles, managing context, and providing human-in-the-loop control. The automation layer executes IT practices in a repeatable and measurable way—generating artifacts, code, documentation, and tests that align with your organization's standards.

Organizations that implement an AI Delivery Operating System gain a capability that no point tool can provide: systemic, scalable AI adoption—from a single pilot team to an entire organization, with consistency, quality, and control at every stage.

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

PART II: ARCHITECTURE AND INTEGRATION WITH RED HAT OPENSHIFT

PractIQ as AI Delivery Operating System

Inteca's PractIQ is the first enterprise-class platform to fully implement the AI Delivery Operating System concept. Designed based on Inteca's many years of experience in SDLC governance implementation projects for the financial, banking, insurance and industrial sectors, PractIQ provides a ready-to-use operating environment for IT departments ready for systemic AI adoption.

The platform is based on the concept of practice – understood not as a rigid sequence of process steps, but as a repeatable, result-oriented way of acting that can be implemented by humans, AI agents, or a combination of these. This approach – practice-first engineering – is the foundation of PractIQ's philosophy and sets it apart from task-automation-oriented platforms.

PractIQ offers two complementary operating models. The Core model provides ready-to-use, proven IT practices tailored to specific technologies and frameworks – an organization can start working immediately, without having to design processes from scratch. Model Forge allows you to design and automate your own practices, tailored to its specifics, internal standards and tool ecosystem. Both models can work in parallel, enabling incremental deployment from out-of-the-box practices to a fully personalized experience.

The platform provides five ready-made SDLC practices, which together form an integrated, end-to-end manufacturing process automation system.

AI Analysis is the practice of AI-assisted analysis of business and technical requirements. AI agents structure, validate, and complete requirements for the knowledge base and industry patterns, ensuring completeness and consistency at the start of a project. Practice eliminates incomplete or conflicting requirements—one of the most common sources of costly errors in the later stages of SDLC.

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

AI Architecture standardizes architectural design in a complex architectural landscape with the support of AI agents. It automatically generates architectural documentation in C4 Model standards, reducing the time to create architecture models by more than 90% while maintaining consistency and compliance with organizational standards.

AI Application Design automates the creation and validation of technical and functional documentation for a single component. It transforms documentation from a one-off, manual task into a repeatable, scalable practice. The produced models are prepared for further use by AI development and testing agents — becoming part of the organization's living knowledge base.

AI Development standardizes and controls the agentic coding process. It defines agent roles, operating rules, and workflows for engineers—eliminating randomness and ensuring consistent code quality regardless of project or team. The practice supports a wide range of technologies: Angular, Spring Boot, React, .NET and others.

AI Quality Assurance embeds continuous quality assurance throughout SDLC. Automates the creation and execution of E2E validation tests for entire user journeys and integrations between components—with full traceability to requirements.

All practices operate in a human-in-the-loop model: AI agents perform tasks, but a human retains full control and the ability to intervene at every stage. This is a key element of the PractIQ architecture — AI works according to standards, does not improvise.

PractIQ Technical Architecture

PractIQ is a modular five-tier system designed for enterprise environments – with full control, traceability, and the ability to adapt to an organization's existing IT ecosystem.

The Human-in-the-Loop Layer provides continuous human control over the actions of AI agents. The layer provides conversational interfaces accessible through a web browser and integration with Visual Studio Code, allowing

engineers to interact with the system naturally when working with the products of the practice. A key function of this layer is the ability to accept, modify, or undo automated actions proposed by agents—in a supervised mode that ensures that a human remains responsible for the final shape of the artifacts.

The Knowledge Layer collects, organizes, and shares the domain and technical knowledge necessary to implement practices. The layer contains practice definitions, rules of conduct, and knowledge items—templates, architectural patterns, organizational standards, and examples. Thanks to this layer, AI agents do not operate on the general knowledge of the language model, but on knowledge tailored to the realities of a particular organization. The layer allows both the use of practices provided by Inteca and the building of one's own repository of organizational knowledge.

The Data Context Layer provides complete organizational context for AI agents through integrations with key IT systems. The layer integrates with JIRA to read and modify the backlog, tasks, and errors; with Enterprise Architect to access architecture diagrams and domain models; and documentation tools such as Confluence or Notion. Based on the data from these systems, the layer builds a knowledge graph that enables agents to make decisions that align with the actual state of the project and organization.

The LLM Execution Layer is responsible for performing tasks based on generative AI. The layer supports multiple language models—OpenAI, GPT, Anthropic Claude, Gemini, and local models—and dynamically selects the model for the task. The layer architecture allows for fine-tuned models on organizational data using LoRA adapters, which optimizes costs while maintaining the quality of results and ensures data sovereignty in environments with increased security requirements.

The Orchestration Layer is the central component that integrates the other layers and manages the flow of data and tasks within the practices. The layer is responsible for configuring and building tool servers (MCP Servers), orchestrating the flow between layers, managing the memory of agent conversations, extracting knowledge from domain systems (Knowledge Ingestion Pipeline), and monitoring and logging all activities. The orchestration layer is also the entry point

for future Agent-to-Agent (A2A) integration—a protocol that enables collaboration between agents from different platforms and organizations.

Integration with Red Hat OpenShift

Red Hat OpenShift is the recommended production environment for PractIQ – providing the enterprise-grade security, scalability, and production readiness required by large enterprise-grade organizations.

Why OpenShift. Implementing PractIQ in an enterprise environment requires a platform that meets enterprise requirements for security, compliance, access management, and scalability – without having to build this infrastructure from scratch. As the leading enterprise-class Kubernetes platform, Red Hat OpenShift delivers all of these elements in a single, integrated environment. For organizations that are already using OpenShift as an application platform, implementing PractIQ means expanding their existing ecosystem with an AI governance and SDLC automation layer – without the need to deploy new infrastructure.

How PractIQ works on OpenShift. All PractIQ components are containerized and deployed as OpenShift workloads, managed by native platform mechanisms. The integration includes a range of Red Hat products, which together form a complete runtime environment for PractIQ.

Red Hat OpenShift is the foundation of the platform, providing container orchestration, autoscaling, environment isolation, and built-in security and compliance. The Orchestration Layer with MCP servers acts as a set of microservices managed by OpenShift, using built-in service discovery, load balancing, and auto-restart mechanisms.

Red Hat OpenShift AI is used to run and manage language models directly in the cluster - including local models and fine-tuned models with LoRA adapters). This ensures full data sovereignty and eliminates the need to send sensitive organizational data to external APIs. OpenShift AI also provides an environment for fine-tuning models on organizational data, which optimizes quality and costs with heavy use.

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

Red Hat OpenShift Dev Spaces provides a development environment integrated into the platform—allowing engineers to work with PractIQ AI agents directly in the browser, without the need to configure environments locally. This is especially important in organizations with distributed teams or security requirements that preclude working on local machines.

The Red Hat Developer Hub serves as a central point of access to PractIQ practices and tools for engineering teams—integrating with your organization's existing developer portal and providing a consistent user experience regardless of project or team.

Red Hat Developer Hub Workflow allows you to run PractIQ practices in a serverless model—specifically for agent orchestration and event processing tasks that require flexible scaling from zero to multiple instances depending on your workload. This reduces operating costs in the event of irregular or seasonal use of the platform.

The Human-in-the-Loop Layer is accessible via OpenShift Routes with built-in TLS, providing secure access for end users without additional reverse proxy configuration.

Benefits of OpenShift integration. Automatic scaling of PractIQ components depending on the workload ensures stable operation even with intensive use by multiple teams at the same time. Isolation of development, test and production environments in accordance with corporate governance standards eliminates the risk of unintentional modifications to production environments. Built-in RBAC, network policies, and containerized image scanning provide enterprise-compliant security without additional effort. Integration with an existing CI/CD pipeline based on OpenShift Pipelines or Tekton allows you to automatically deploy PractIQ updates as part of your organization's standard DevOps processes.

Deployment models. PractIQ on OpenShift can be deployed in three configurations. On-premise deployment in an OpenShift cluster provides complete control over infrastructure and data – optimal for organizations with the highest security requirements, such as banks and financial institutions. A

hybrid cloud deployment combines an on-premises OpenShift cluster for sensitive components with cloud resources for components that require elastic scale. Multi-cloud deployment on managed OpenShift services (ROSA on AWS, ARO on Azure) enables quick commissioning without investing in your own infrastructure.

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

PART III: USE CASES AND BUSINESS VALUE

Application scenarios

PractIQ with Red Hat OpenShift addresses the full spectrum of AI-driven delivery challenges, from new projects to existing systems to legacy upgrades. The three scenarios correspond to three real-world situations faced by IT departments of large enterprise organizations.

Scenario 1: Forward Engineering – New Projects (Greenfield)

Organizations building new systems or developing existing digital products face a fundamental challenge: how do you make every new project use AI in a consistent, controlled, and repeatable way—regardless of team, technology, or location?

Without a structured operating model, every project starts from scratch. Engineers experiment with AI agents on their own, the results are unpredictable, and the quality of the code is variable. IT works faster, but not better.

PractIQ provides a ready-made, AI-powered operating model for full SDLC. The five automated practices — Analysis, Design, Architecture, Development, QA — operate as an integrated system where each practice uses artifacts produced by the previous one. AI agents perform tasks according to defined standards, and engineers retain full control in a human-in-the-loop model.

Effects measured in Inteca projects: time-to-market improved by 40%, 60% increase in team efficiency, over 90% compliance with governance, 95% code generated by AI under the supervision of engineers.

Scenario 2: Backward Engineering - Existing Systems (Brownfield)

Imagine an IT department managing hundreds of applications—most of which are running, but no one can tell exactly how. Every attempt at modernization starts with the same question: how do you know what can be changed without breaking what works? When AI agents come to this picture—who need precise context to generate actionable results—the question becomes urgent.

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

PractIQ solves this problem by automatically restoring missing documentation and architecture directly from the source code. AI agents analyze existing repositories, identify the domain structure, map dependencies between systems, and generate up-to-date documentation in the C4 Model and ADR standards. The result goes to the PractIQ Knowledge Layer — creating an AI-ready knowledge base that becomes the foundation for further work of agents.

An organization is moving from a state where no one fully understands what applications are doing to a state where every agent—and every new engineer—has access to up-to-date, structured knowledge of systems.

Scenario 3: Modernization and elimination of vendor lock-in (Legacy & Modernization)

Once your organization understands its systems—with the reconstructed documentation and architecture from Scenario 2—PractIQ enables the next step: proactive modernization. Rewriting applications in new technologies, replacing systems controlled by external vendors, reducing technological debt – all this without the risk of losing domain knowledge.

The traditional approach to modernization is fraught with high risk: teams rewrite systems they don't fully understand, leading to missing key business logic and costly errors. PractIQ eliminates this risk – AI agents work on the basis of the knowledge recreated in the previous stage, generating new code that complies with the business domain and standards of the organization.

The three key benefits of this scenario are modernization with minimal risk (rewriting applications in new technologies without losing domain knowledge), reducing technological debt (replacing outdated architecture with a clean, AI-driven foundation), and eliminating vendor lock-in (restoring and taking control of systems managed by external vendors).

Three scenarios create a natural path to maturation: organizations can enter anywhere—starting with new projects, organizing existing systems, or urgently modernizing—and move at their own pace and priorities.

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

Business Value

Implementing PractIQ with Red Hat OpenShift delivers measurable business value on four levels – with the specific benefits varying depending on the scenario in which the organization begins its journey.

The operational results achieved by Inteca in its own manufacturing process and in the first customer projects show the scale of possible improvements in the Greenfield scenario. The workload for creating functionalities is reduced by up to 80%. The analysis and architecture phase is shortened by more than 50%. The bandwidth and speed of teams increases by 50%. The platform provides 90% test coverage without additional engineer effort. Time-to-market improves by 40%. 95% of the code is produced by AI agents under the supervision of engineers, and 80-90% of the technical specification is created automatically.

The value from the Brownfield and Modernization scenarios is of a different kind, but equally measurable. Organizations that reconstruct the documentation and architecture of existing systems benefit primarily from a reduction in operational risk – architectural and modernization decisions are made based on current, structured knowledge, rather than guesswork and residual institutional memory. The onboarding time of new engineers is shortened, and instead of weeks spent on reconstructing knowledge about systems, they receive a ready-made database of documentation. Eliminating vendor lock-in brings direct financial benefits – independence from external vendors, regaining control over licensing and maintenance costs, and the ability to modernize on your own terms. According to Gartner's research, technology debt consumes an average of 30-40% of the IT budget in large organizations — PractIQ provides a systematic tool for its reduction.

The financial benefits common to all scenarios result from the automation of repetitive IT practices and the elimination of the cost of reproducible work. Reducing rework and rework—resulting from higher quality and consistency of artifacts produced—reduces the cost of late project phases, where errors are most costly. The LLM architecture based on dynamic model selection and fine-tuning capabilities with LoRA optimizes token costs. For organizations that want to

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

maintain full data sovereignty, deploying on-premises models on OpenShift AI eliminates variable API costs with heavy usage.

The strategic value of PractIQ is primarily due to the fact that all three scenarios run on a single, integrated platform. An organization doesn't have to choose between investing in new projects or tidying up existing systems—PractIQ supports both simultaneously, and the knowledge generated in the Brownfield scenario becomes a direct contribution to the Greenfield and Modernization scenarios. Governance and compliance built into every workflow eliminate the operational risk associated with the uncontrolled use of AI. Independence from specific AI model vendors reduces vendor lock-in risk at the platform level. Deploying on Red Hat OpenShift ensures that PractIQ fits naturally into your organization's existing platform strategy, minimizing deployment risk and integration costs.

Deployment model

PractIQ is designed for incremental deployment – an organization can start with a single scenario and scale at its own pace, avoiding the risk of a large big-bang deployment.

Stage 1: Preparations. An organization prepares to work with AI by organizing architecture, documentation, and data. PractIQ AI Architecture and AI Documentation allow you to quickly build an AI-ready knowledge base—whether you're starting from new projects or existing systems that require documentation reproduction.

Stage 2: Standardization. The organization implements the first SDLC practices and standardizes the way of working with AI agents. Engineers learn to work in a human-in-the-loop model, and the IT department develops internal standards and rules for agents.

Stage 3: Automate. The organization implements a full set of SDLC practices and automates the end-to-end manufacturing process. The AI Software Factory becomes an operational reality: AI agents carry out the entire production cycle under the supervision of engineers.

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

Stage 4: Scale. The organization is expanding PractIQ to more departments, projects, technologies, and business areas — including legacy modernization and elimination of vendor lock-in. The Forge model allows you to create your own practices tailored to the specifics of new areas.

Each stage has a clearly defined entry, exit, and business impact — allowing you to measure progress and justify subsequent investments to the organization's board.

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

Summary

The era of AI is changing software development as profoundly as agile and DevOps once did. But unlike previous transformations, this change doesn't just affect tools and processes — it touches on a fundamental question about how humans and machines work together in software development.

At the same time, organizations are entering this transformation with the baggage that the previous era — the Agile era — left them behind: hundreds of applications with architectural knowledge dispersed and uncodified, with no up-to-date documentation, with mounting technological debt. Systemic AI adoption requires facing both challenges at the same time — you can't build the future on a foundation that no one fully understands.

IT departments of large enterprise organizations that try to respond to these challenges exclusively with point tools are facing a wall: fragmented experiments, unpredictable results, lack of scalability, and the increasing risk of losing control of the manufacturing process.

The AI Delivery Operating System is the answer to this gap—not as another tool, but as a new operational layer that organizes the way IT works in the AI era. PractIQ puts this concept into practice, providing a complete environment for three scenarios: building new systems with AI from scratch (Greenfield), recreating knowledge about existing applications and preparing them to work with agents (Brownfield), and modernizing and eliminating vendor lock-in legacy systems (Modernization). All three scenarios run on a single, integrated platform, and the knowledge generated in each scenario feeds the others.

Red Hat OpenShift provides the enterprise-grade foundation for this transformation—the security, scalability, and production readiness required by organizations operating in regulated industries and complex IT environments.

Together, PractIQ and Red Hat OpenShift create a complete solution for IT departments ready for the next step. Because AI doesn't scale by adding more tools. It scales by building an operating model that manages how AI people and agents work together — in new projects, existing systems, and legacy upgrades.

Inteca Sp. z o.o.
Powstańców Śląskich 9
53-332 Wrocław

+48 71 715 60 91
biuro@inteca.pl

inteca.com
practiq.tech

About companies

Inteca is an IT company based in Wroclaw, founded in 2011, employing over 60 engineers and consultants. Inteca specializes in the design, implementation and maintenance of enterprise-class IT platforms, providing services in the areas of Identity and Access Management, Business Automation, Cloud, Data Integration, DevOps and Enterprise Architecture. The company serves clients in the financial, banking, insurance, industrial and energy sectors — including Santander, BNP Paribas, Credit Agricole, Allianz, Generali, Orlen and Tauron.

Inteca is a Red Hat Advanced Business Partner with many years of experience in consulting, deployment, and managed services for Red Hat OpenShift and Kubernetes. PractIQ is an extension of this specialization — as an AI governance and SDLC automation platform complementary to the Red Hat ecosystem.

For more information: inteca.com | practiq.tech

Red Hat is the world's leading provider of open source solutions for enterprises, offering proven technologies for Linux, cloud, container, and Kubernetes environments. Red Hat OpenShift is the leading enterprise-class Kubernetes platform, providing organizations with a secure, scalable environment to run containerized applications in on-premise, hybrid cloud, and multi-cloud environments. Red Hat helps customers integrate new and existing IT technologies, develop cloud applications, and automate and manage complex environments.

More information: redhat.com

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