

Red Hat saves \$5 million with AI-powered innovation in IT support



Headquarters

Raleigh, North Carolina, U.S.

Industry

Technology

Size


1,100 global support associates and engineers

"AI augmentation doesn't just improve efficiency; it also enhances content creation and may contribute to job satisfaction. This will be crucial as Red Hat scales to handle increasing case volumes without overburdening associates."

Mandy Elliott

Senior Director
AI and Data
Experience Engineering
Red Hat

Red Hat sought to increase the efficiency and scalability of customer and technical support services to its growing customer base with AI solutions. The Experience Engineering (XE) team at Red Hat embarked on a program, using the Red Hat AI portfolio comprised of Red Hat OpenShift AI and Red Hat Enterprise Linux AI, to develop, test, and deploy 4x solutions powered by artificial intelligence (AI), all with the aim of simplifying IT support for customers and support associates. These tools improve self-service, increase efficiency, and help bring about a faster response to support cases. As well as delivering cost avoidance in the region of \$5 million, the AI solutions enhance job satisfaction for IT support associates while improving customer experiences.



Software and services

Red Hat® OpenShift® AI
Red Hat Enterprise Linux® AI

Partner

IBM

Benefits

- ▶ Delivered over \$5m in cost avoidance, with estimated \$1.5m in just 10 months
- ▶ Increased availability of knowledge content and minimized repetitive tasks for IT support associates handling 30,000 new cases each month
- ▶ Provided faster responses to customers with AI, enhancing overall experience

About Red Hat Innovators in the Open

Innovation is the core of open source. Red Hat customers use open source technologies to change not only their own organizations, but also entire industries and markets. Red Hat Innovators in the Open proudly showcases how our customers use enterprise open source solutions to solve their toughest business challenges. Want to share your story? [Learn more.](#)

“The troubleshooting optimization project alone has avoided an estimated \$1.5 million in support costs in just 10 months, by reducing the number of users who open a case after clicking on our content.”

Mike Clark

Senior Manager
Software Engineering
Red Hat

Increasing support efficiency to align with business growth

Red Hat’s support services are a cornerstone of the company’s success and a key factor in why customers choose its technologies. With a growing customer base, case volumes are at an all-time high and continue to rise.

Red Hat’s global team of 1,100 technical support associates and engineers operates 24x7, handling approximately 30,000 new cases each month. To ensure support is scalable and efficient, the Red Hat Experience Engineering (XE) team within Red Hat Global Engineering was tasked with making support easier for associates while enhancing customer experience with artificial intelligence.

The team aimed to minimize repetitive tasks and reduce the time spent searching for the right information to bring about faster responses and more efficient problem resolution. “With vast amounts of support data available, we wanted to maximize the use of AI and make it more accessible while ensuring data privacy,” said Mandy Elliott, Senior Director, AI and Data, Red Hat Experience Engineering.

Taking advantage of AI to augment IT support

The team launched a new innovative AI-driven program designed to embrace flexibility and an experimental mindset to ensure success. They created a ‘betmap’ – an iterative experiment-based alternative to a traditional roadmap. “For our initial foray into AI for support, we selected use cases with structured or semi-structured data to maximize success, such as knowledge-centered service (KCS) solution content, which already use industry-standard formats and avoid privacy concerns,” said Manikandan Sivanesan, AI Technical Strategy Lead, Red Hat Experience Engineering.

The program focused on 4 key projects using the Red Hat AI portfolio: enhancing the Red Hat Customer Portal Case Management Troubleshooting tool, creating an AI summarization model to generate concise summaries of solutions, developing an AI-powered summarization tool to support case handovers, and an AI-powered tool to automate the creation of standardized knowledge base articles.

The team initially experimented with AI models on Red Hat Enterprise Linux AI, which supports the development of Granite family large language models (LLMs) for enterprise applications. Validated applications were then scaled and deployed on Red Hat OpenShift AI on AWS, an enterprise-ready AI and machine learning (ML) platform designed for hybrid cloud environments.

The first project enhanced the Red Hat Customer Portal Case Management Troubleshooting tool, optimizing the order of knowledge articles provided to users based on relevance to issue descriptions. The XE team, working in partnership with IBM Research and Red Hat IT, developed an AI-powered textual reranking solution, achieved using a fine-tuned IBM Slate Retriever model running on Red Hat IT’s OpenShift AI cluster, with GPU acceleration.

With approximately 130,000 knowledge base solutions, many previously undiscoverable, the team created an AI summarization model to generate accurate knowledge base snippets. Running on IBM watsonx.ai and OpenShift AI’s data science pipelines, the model generates concise summaries of solutions.

High-severity support cases often require handovers across associates, time zones, and teams. To streamline these transitions, the team developed an AI-powered summarization tool that generates concise summaries of support case interactions on demand. Initially designed for 24x7 'follow the sun' cases, it has now been expanded to additional support groups. Built using Mistral LLMs on Red Hat IT's OpenShift AI cluster with GPU acceleration, the team is now testing Granite models and exploring further fine-tuning with InstructLab – an open-source project by IBM and Red Hat.

The team also piloted an AI-powered tool to automate the creation of standardized knowledge base articles, aligning with KCS v6 standards: "Our tool uses support case comments to draft articles quickly and consistently, reducing reliance on manual efforts," said Mike Clark, Senior Manager, Software Engineering, Red Hat Experience Engineering. "Built using Mistral LLMs and GPU acceleration on OpenShift AI, this solution enhances self-service capabilities, allowing support associates to focus on solving new issues."

Delivering cost avoidance while improving the customer and IT associate experience

Delivered over \$5m in cost avoidance, with estimated \$1.5m in just 10 months

Red Hat's AI-powered initiatives are projected to save more than \$5 million through successful customer self-serve and reduction of manual effort. Cost avoidance means the organization is able to address growing demand without increasing its budget. "The troubleshooting optimization project alone has saved an estimated \$1.5 million in support costs in just 10 months, by reducing the number of users who open a case after clicking on our content," said Clark.

"Additionally, automating the drafting of KCS articles is expected to generate an annual cost avoidance impact of \$8 million."

Increased availability of knowledge content and minimized repetitive tasks for IT support associates

AI solutions help IT support associates work more efficiently by providing faster access to relevant information and automating repetitive tasks. "AI-assisted solutions can streamline self-service processes and help support associates respond more quickly," said Elliott.

This then frees support associates to focus on more meaningful and rewarding activities. "AI augmentation doesn't just improve efficiency – it also enhances content creation and may contribute to job satisfaction," said Elliott. "This will be crucial as Red Hat scales to handle increasing case volumes without overburdening associates."

Provided faster responses to customers with AI, enhancing the overall experience

Knowledge articles can now be provided to users in order of relevance based on issue descriptions, which in turn speeds up customer access to articles. Moreover, the AI summarization model can now generate knowledge base snippets and concise summaries of solutions to support users in quickly identifying relevant information.

"We've streamlined the user experience by making content scannable and easier to find," said Clark.

Empowering customers to self-serve effectively improves their overall experience and reduces resolution times. "AI not only makes self-service more accessible but also ensures faster, more accurate responses when customers need support from Red Hat," said Sivanesan. "With support managed consistently, our customers can stay focused on serving their own clients."

Establishing best practices in AI development

Developing these solutions has enhanced Red Hat's IT support capabilities and established best practices for AI development projects. Key lessons include the importance of maintaining project focus and leveraging multidisciplinary expertise when needed. "With such vast AI capabilities, it's easy for projects to spiral out of control," said Sivanesan. "While flexibility and an experimental mindset are essential, it's equally important to stick to the project scope."

Red Hat's solutions provided the XE team with the extensive feature set and flexibility required for success. "With Red Hat, you're not locked into a single approach. We used a variety of techniques and models, selecting what worked best for each situation," said Sivanesan.

Building on the success of these initial projects, the team is now transitioning pilots into enterprise-wide programs, deploying them in production environments, and exploring opportunities to scale and replicate its AI-powered innovations.

About Red Hat Experience Engineering (XE)

Red Hat's global team of 1,100 technical support associates and engineers operates 24x7, handling 30,000 new cases each month. XE also produces the [Red Hat Customer Portal](#) and official [Red Hat Product Documentation](#) to support customers in their journey with Red Hat.



About Red Hat

Red Hat is the world's leading provider of enterprise open source software solutions, using a community-powered approach to deliver reliable and high-performing Linux, hybrid cloud, container, and Kubernetes technologies. Red Hat helps customers develop cloud-native applications, integrate existing and new IT applications, and automate and manage complex environments. [A trusted adviser to the Fortune 500](#), Red Hat provides [award-winning](#) support, training, and consulting services that bring the benefits of open innovation to any industry. Red Hat is a connective hub in a global network of enterprises, partners, and communities, helping organizations grow, transform, and prepare for the digital future.

North America

1 888 REDHAT1
www.redhat.com

Europe, Middle East, and Africa

00800 7334 2835
europe@redhat.com

Asia Pacific

+65 6490 4200
apac@redhat.com

Latin America

+54 11 4329 7300
info-latam@redhat.com