

Jet Brains



JetBrains Centralizes Developer Support at Scale With Rasa

How a global leader in intelligent software development tools uses Rasa to unify a fragmented support ecosystem and give millions of developers a faster path to answers

75–80%

annual customers served

20+

products supported

80%

deflection rate

JetBrains creates intelligent software development tools trusted by over 15 million users and 88 Fortune Global Top 100 companies. Headquartered in Amsterdam, the company operates as a key enabler of technology, supporting the developers who build the industry-shaping products of the future. Its lineup of 30+ products includes award-winning IDEs like IntelliJ IDEA, the AI coding agent Junie, and Mellum - JetBrains' focal LLM purpose-built for code-related tasks. These milestones, alongside team tools like YouTrack and the multiplatform language Kotlin, exemplify a commitment to reducing development friction by natively embedding emerging technologies into professional workflows.

INDUSTRY

Software

LOCATION

Amsterdam,
The Netherlands

EMPLOYEES

2,800

SCALE OF OPERATIONS

Used by over 12.8 million professionals and 92 of the Fortune Global Top 100

Key Takeaways

- **Centralized entry point:** Rasa powers a single AI-driven support agent that routes customers to the right JetBrains team instantly, replacing a fragmented web of 100+ separate contact forms.
- **High customer satisfaction:** JetBrains is sustaining a CSAT score of 75–80% with their Rasa-powered agent, which is approaching parity with their human support team.
- **Scalable without compromise:** JetBrains' Rasa-powered agent deflects a rapidly scaling volume of routine inquiries so humans can focus on the complex, highly technical cases they're best equipped for.
- **On-premise control:** Rasa's self-hosted deployment model gave JetBrains full control over sensitive customer data, eliminating legal and compliance delays and enabling a faster launch.

The Challenge

JetBrains customers are developers with highly technical support needs that require precise answers. Building a specialized support operation capable of meeting their expectations across a portfolio of 20+ products was itself a considerable undertaking. Doing it at scale for millions of users made it even harder.

A SUPPORT MODEL BUILT FOR QUALITY, NOT SCALE

For most of JetBrains' history, their support model was built around depth, not throughput. Each ticket received careful attention from technically skilled agents, many of whom specialized in specific areas of the IntelliJ platform.

The result was consistently high satisfaction. But, as JetBrains' customer base grew, the cracks in the model became harder to ignore:

- Resolution times stretched.
- Tickets required multiple back-and-forth exchanges to gather enough context for a resolution.
- Agents found themselves repeatedly answering the same known issues.

It was a quality-first model straining under the weight of volume, and it was occupying human expertise that could be better directed at genuinely novel problems.

"We realized that this particular approach was not scalable," said Dino Letic, project manager at JetBrains. "We had a lot of users waiting for their tickets to be answered, and long resolution times were common."

A FRAGMENTED ENTRY POINT EXPERIENCE

The scalability problem was compounded by a structural one. Over time, JetBrains had accumulated more than 100 separate contact forms and support entry points scattered across their website and products. Customers who wanted to get help had to first figure out where to submit their inquiry.

"It shouldn't be difficult for customers to figure out how to reach JetBrains," Letic noted. "We wanted to give customers one entry point to make their support experience as easy as possible."

"It shouldn't be difficult for customers to figure out how to reach JetBrains. We wanted to give customers one entry point to make their support experience as easy as possible."

DINO LETIC
Project Manager
at JetBrains

This fragmentation also made it difficult to operate efficiently. Running analytics across disconnected intake channels was unwieldy. And identifying patterns in customer issues — the kind of insights that drive meaningful product and support improvements — required piecing together data from too many places.

The Solution

JetBrains evaluated several vendors before selecting Rasa. Two requirements ultimately drove their decision: on-premise deployment capability and a platform that could serve as the foundation for a long-term, evolving support operation.

ON-PREMISE DEPLOYMENT: A NON-NEGOTIABLE

As a company that handles sensitive customer data at scale, routing that data through a third-party cloud wasn't an option for JetBrains. Rasa's ability to run fully self-hosted on JetBrains' own infrastructure was, in Letic's words, "one of the main reasons" they chose it. The fact that it would be on-prem also removed the need to negotiate data processing agreements and push through lengthy legal reviews, which let JetBrains move from evaluation to launch faster than a cloud-only alternative would have allowed.

BUILDING THE AGENT

The JetBrains team used Rasa Studio to design and develop conversation flows, giving team leads visibility into the agent's logic without requiring deep engineering involvement at every step.

The agent they built does several things well:

- It recognizes customer intent across the full breadth of JetBrains' product catalog.
- It determines whether a question can be resolved immediately or needs escalation.
- When it can answer, it generates a response grounded in JetBrains' internal knowledge base, structured through their KCS (Knowledge-Centered Service) framework.
- When it can't, it escalates to the appropriate human team with full context already captured, so customers don't have to explain their issue again.

Deployed

on-premise infrastructure to meet enterprise data compliance requirements without lengthy legal review

Unified

intent recognition across JetBrains' full product catalog to route, resolve, or escalate with full customer context preserved

Eliminated

the need for customers to repeat themselves by passing full conversation context to human agents at the point of escalation

JetBrains also implemented pre- and post-processing checks to validate that an incoming question is complete before generating a response and to verify that the response itself meets quality standards before it reaches the customer.

"We wanted something that's not just a deflection tool," Letic said. "Rasa allowed us to build an advanced chatbot that can understand complex customer needs and problems."

The Results

JetBrains' Rasa-powered agent is now live across the key entry points of their support ecosystem, processing roughly 3,000 conversations per month and sustaining a CSAT score of 75–80% — numbers that put the agent on par with their human support team in terms of customer satisfaction.

Customer feedback suggests users are happy when they get answers instantly, and they appreciate that, when escalation happens, it's a smooth experience. Many conversations end with a simple "thank you" and no further action — a signal that the agent is resolving their issues.

The centralized entry point is also delivering on its operational promise. Instead of tracking outcomes across fragmented channels, JetBrains now has a single system where patterns are visible, analytics are actionable, and improvements can be measured clearly.

The Future

JetBrains is actively expanding the agent's coverage to additional entry points across their web presence. They're also developing more sophisticated use cases in sales-sensitive areas like license management and account issues. Longer term, they expect their Rasa agent will handle the majority of JetBrains' support conversation volume.

"We wanted something that's not just a deflection tool. Rasa allowed us to build an advanced chatbot that can understand complex customer needs and problems."

DINO LETIC
Project Manager
at JetBrains