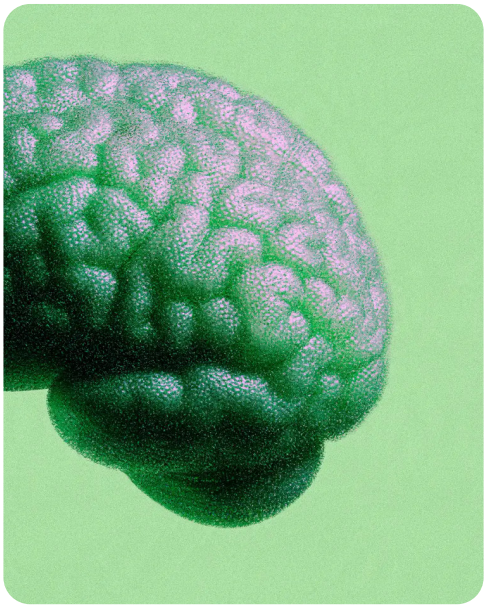
A Slack knowledge assistant for Nezasa, built in one week

A one-week build on Nezasa's existing data platform put an assistant in Slack that answered 55% of evaluated questions without a product expert.



CLIENT

Nezasa AG

INDUSTRY

Travel Tech

PUBLISHED

April 2024

READ ONLINE

samuylov.ai · A Slack knowledge assistant for Nezasa, built in one week

SERVICES

GenAI opportunity framing · Retrieval-augmented generation · Data extraction and modelling · Slack integration and rollout

55%
Questions answered without involving a product expert

~2 FTE
Capacity benefit estimated by the client

1 week
From kick-off to a working assistant

Client

Nezasa is a B2B travel technology provider founded in Zürich in 2012, with a team across Switzerland, Portugal and Chile. What its staff needed to sell, support and onboard customers was well documented, but spread across a company wiki, a public knowledge base, API specifications and commercial documents.

Challenge

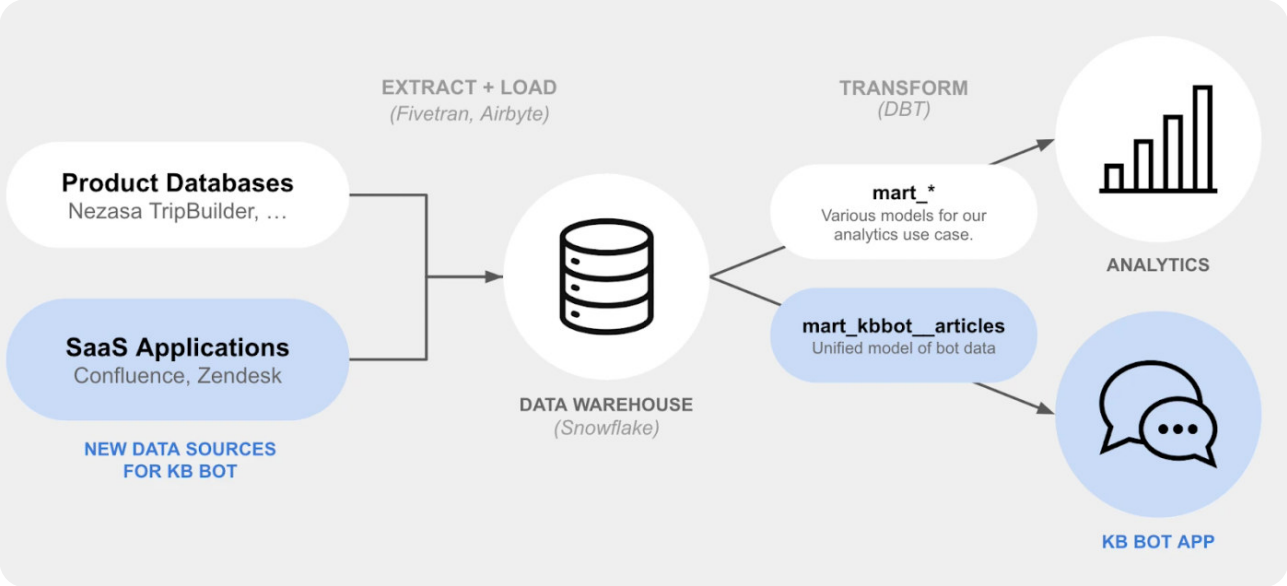
The answers existed; finding them under time pressure was the problem. Sales, support and onboarding questions ended up with the few product and engineering specialists who knew where to look, and every answer cost them time they would otherwise have spent on the product.

The fix had to be small, because Nezasa had little engineering capacity to spare. It also had to draw on every source people relied on, PDFs included, and keep company data under Nezasa's control.

Our approach

Before writing any code, we fixed the goal: show practical value with a build small enough to finish in a week.

Build on the platform that already exists. A one-week build was possible because the assistant extended the data platform from our preceding engagement. Its existing tools were pointed at Confluence and Zendesk to combine their content into one table of articles, and PDFs such as API specifications and commercial terms were added directly.



White: the platform Nezasa already had. Blue: what the assistant added – two data sources, one data model and the assistant itself.

Retrieval, not fine-tuning. A weekly job converts each document to Markdown, splits it at its headings and indexes the sections, keeping the source link with each one. Updating what the assistant knows means re-indexing rather than retraining, and switching model providers stays cheap.

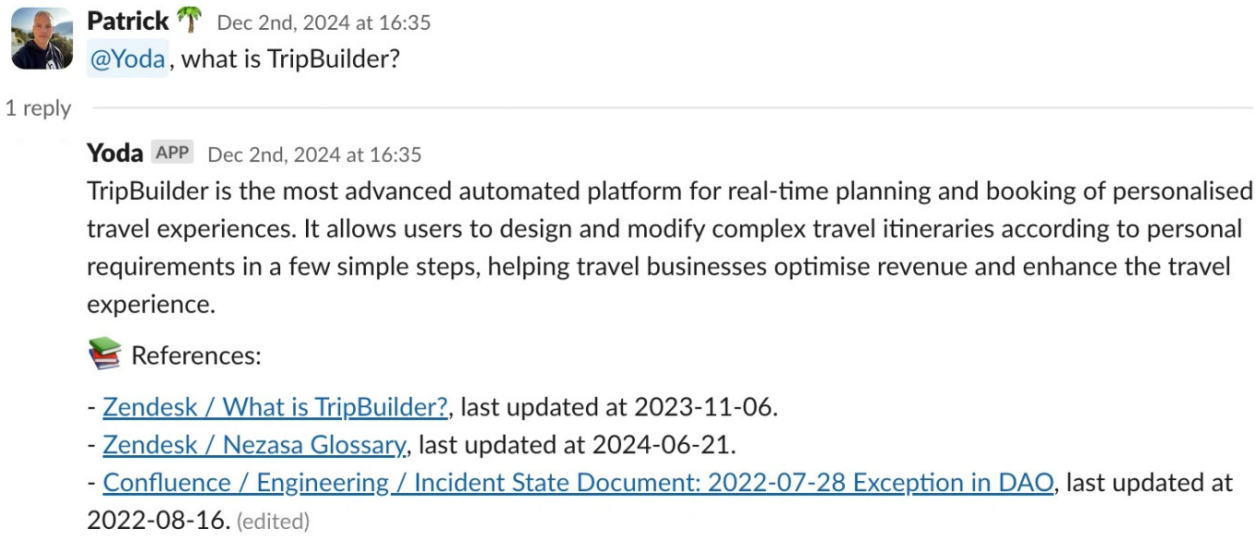
Answer where people already work. Instead of building a new interface, we put the assistant in Slack. People mention it in the channel they are already in and get the answer in the thread, with links to its sources; follow-up questions go in the same thread.

Keep sensitive data out. Only Nezasa employees could use the assistant, and client-owned and sensitive data was excluded from the index from the start.

Results

55% of evaluated questions were answered without involving a product expert. Nezasa estimated the capacity released at roughly two full-time employees. Most of the rest pointed to topics nobody had documented yet – exactly the gaps worth fixing.

People across the company, from the CEO to onboarding managers and UX designers, used it the same way, in Slack, and rated each answer with a thumbs up or down; Nezasa tracked the ratings in a dashboard.



Asked in Slack, answered in the thread, with links to the Zendesk and Confluence pages behind the answer.

Nezasa's own team went on to run and extend the assistant, and later added an API for external partners.

What transfers

Start from the data foundation you already have. Most of the work in a knowledge assistant is getting current content out of the systems where it lives. Where a data platform already does that, the assistant is a short extension rather than a new project.

Meet people in the tool they already use. An assistant inside Slack needs no new tool to learn, and a source link in every answer lets people check it before they rely on it.

Unanswered questions are a to-do list for your documentation. When the assistant misses because a topic was never written down, the miss shows which page is missing. At Nezasa, the unanswered questions showed where the documentation needed work.

Further reading

Denis presented the business case behind this project at GenAI Zürich 2024; it is written up in A one-week knowledge-assistant build at Nezasa. The build itself is documented in a four-part series written with Patrick Hammer:

1. The idea and the business case
2. Design, and the options we rejected
3. Data extraction and preparation
4. Deployment, adoption, and what we would change

Want an assistant that answers from your own documentation? The build can be short. Keeping the answers accurate as documents and questions change is the longer job, and one we can run for you. See Agents-as-a-Service, or Start a conversation.



Denis Samuylov

– Founder & CEO

Founder-led, team-delivered – Denis leads the engagement and stays on it after the pitch.

contact@samuylov.ai +41 79 136 38 26 Book an intro call →