



Build Production-Ready Agentic Workflow

Develop, deploy, and manage agentic workflow, RAG pipelines, and more for teams at any scale, effortlessly.

01

The Challenge

The Transformation Challenge In the Age of AI

48%

AI projects actually
make it into **production**

Technical fragmentation creates infrastructure chaos across databases, models, and deployment system.

70%

AI teams struggle with
data integration

Knowledge is scattered across documents, emails, wikis, and databases.

25%

Employees use AI in their
daily work

Ongoing friction between business and technical teams limits adoption speed and confidence.

1/3

Companies have **trained**
25% of their workforce to use AI

Most organisations are still early in building AI fluency and comfort across teams.

Dify's Solution

1

Production-Ready Since Day One

- Visual workflow builder means no DevOps bottleneck
- Enterprise infrastructure included - security, scaling, monitoring
- Ship to production on day one, iterate in real-time

2

Break Down the Data Silos

- Connect every data source - Notion, Slack, your databases, PDFs, wherever knowledge lives
- RAG pipeline turns fragmented information into intelligent answers
- Update once, reflected everywhere with efficient knowledge management

3

Making AI Accessible for Everyone

- Collaborate to build and iterate workflow smoothly
- Teams with limited technical background can create AI agents in Dify with minimal learning curve
- Integrate with current systems or platforms in minutes

Build with Dify

• BUILD WITH DIFY

Chatbot
Backendless frontend [application](#) development via API integration.

Agent
Conversational [Chatbots](#) to intelligent [Agents](#) and creative [Content Generators](#).

Content Generation
Conversational Chatbots to intelligent Agents and creative Content Generators.

Chatflow
Capable for multi-round complex dialogue tasks with memory.

Workflow
Automate any task using advanced logic like custom Code, IF/ELSE conditions, and Loop.

• BUILD & DESIGN

Dify Studio
Create, debug, and manage AI application in one place.

Built-in Templates
Ready-to-use application templates for common use cases.

Prompt IDE
Visually design, test, and optimize prompts.

RAG Pipeline
Design knowledge data processing procedures for accurate, context-aware retrieval.

• CONNECT & INTEGRATE

Large Language Model
Connect to and switch between global LLM providers.



MCP Server
Tap into a growing ecosystem of MCP servers without complex setup.



Tools and Plugins
Integrate and empower AI application capabilities with external tools.



Data Source
Sync knowledge data from local file, website, online doc, and online drive instantly.



• ORCHESTRATE & TEST

Debugging & Fine-tuning
Test individual nodes, visualize data flow, and debug your AI's logic in real-time to ensure reliability.

Logs & Annotations
Backendless frontend application development via API integration.

• SCALE & MONITOR

Backend as a Service
Fully managed backend that includes APIs, user management, and logging.

Batch Run
Process large datasets or automate repetitive tasks by running workflow on multiple data inputs simultaneously.

Publish as an MCP Server
Turn Dify-built workflow or agent as a standard MCP endpoint that any other MCP client can directly call.

Integrate with Current System
Seamlessly connect to your existing software, databases, and internal workflows.

Embed into Website & Share with WebApp
Put AI into work right away and promote your AI application to the world at ease.

LLMOps
Manage your application's entire lifecycle with built-in or third-party LLMOps tools for continuous improvement.



Security and Compliance Top Priority

Since day one of product design, Dify has made user data security, privacy protection, and compliance a top priority. Dify is SOC 2 Type I & Type II audited, ISO 27001:2022 certified, and fully GDPR compliant, ensuring comprehensive protection of your data security and privacy.



General Data Protection Regulation 



ISO 27001 



SOC2 Type I & Type II 

Community-Driven Globally Recognized

Ranked as one of the top 100 open-source repositories of all time.

☰ APP

1M+

Powered by Dify

☰ LOCAL DEPLOYMENT

550M+

Deployment

🌐 GLOBAL REACH

150+

Countries

👤 CONTRIBUTION

1,000+

Contributors

☆ POPULARITY

110K+

GitHub Stars

🏢 ENTERPRISE READY

60+

Industries

Trusted by Leaders Across Industries Worldwide



Changing the World with **Partners Worldwide**



Strategic Partners

We collaborate with the world's leading technology companies to provide a secure, scalable, and powerful foundation for your AI applications.



Marketplace Partners

We work with a growing ecosystem of partners to build powerful tools and integrations that extend your AI's capabilities.



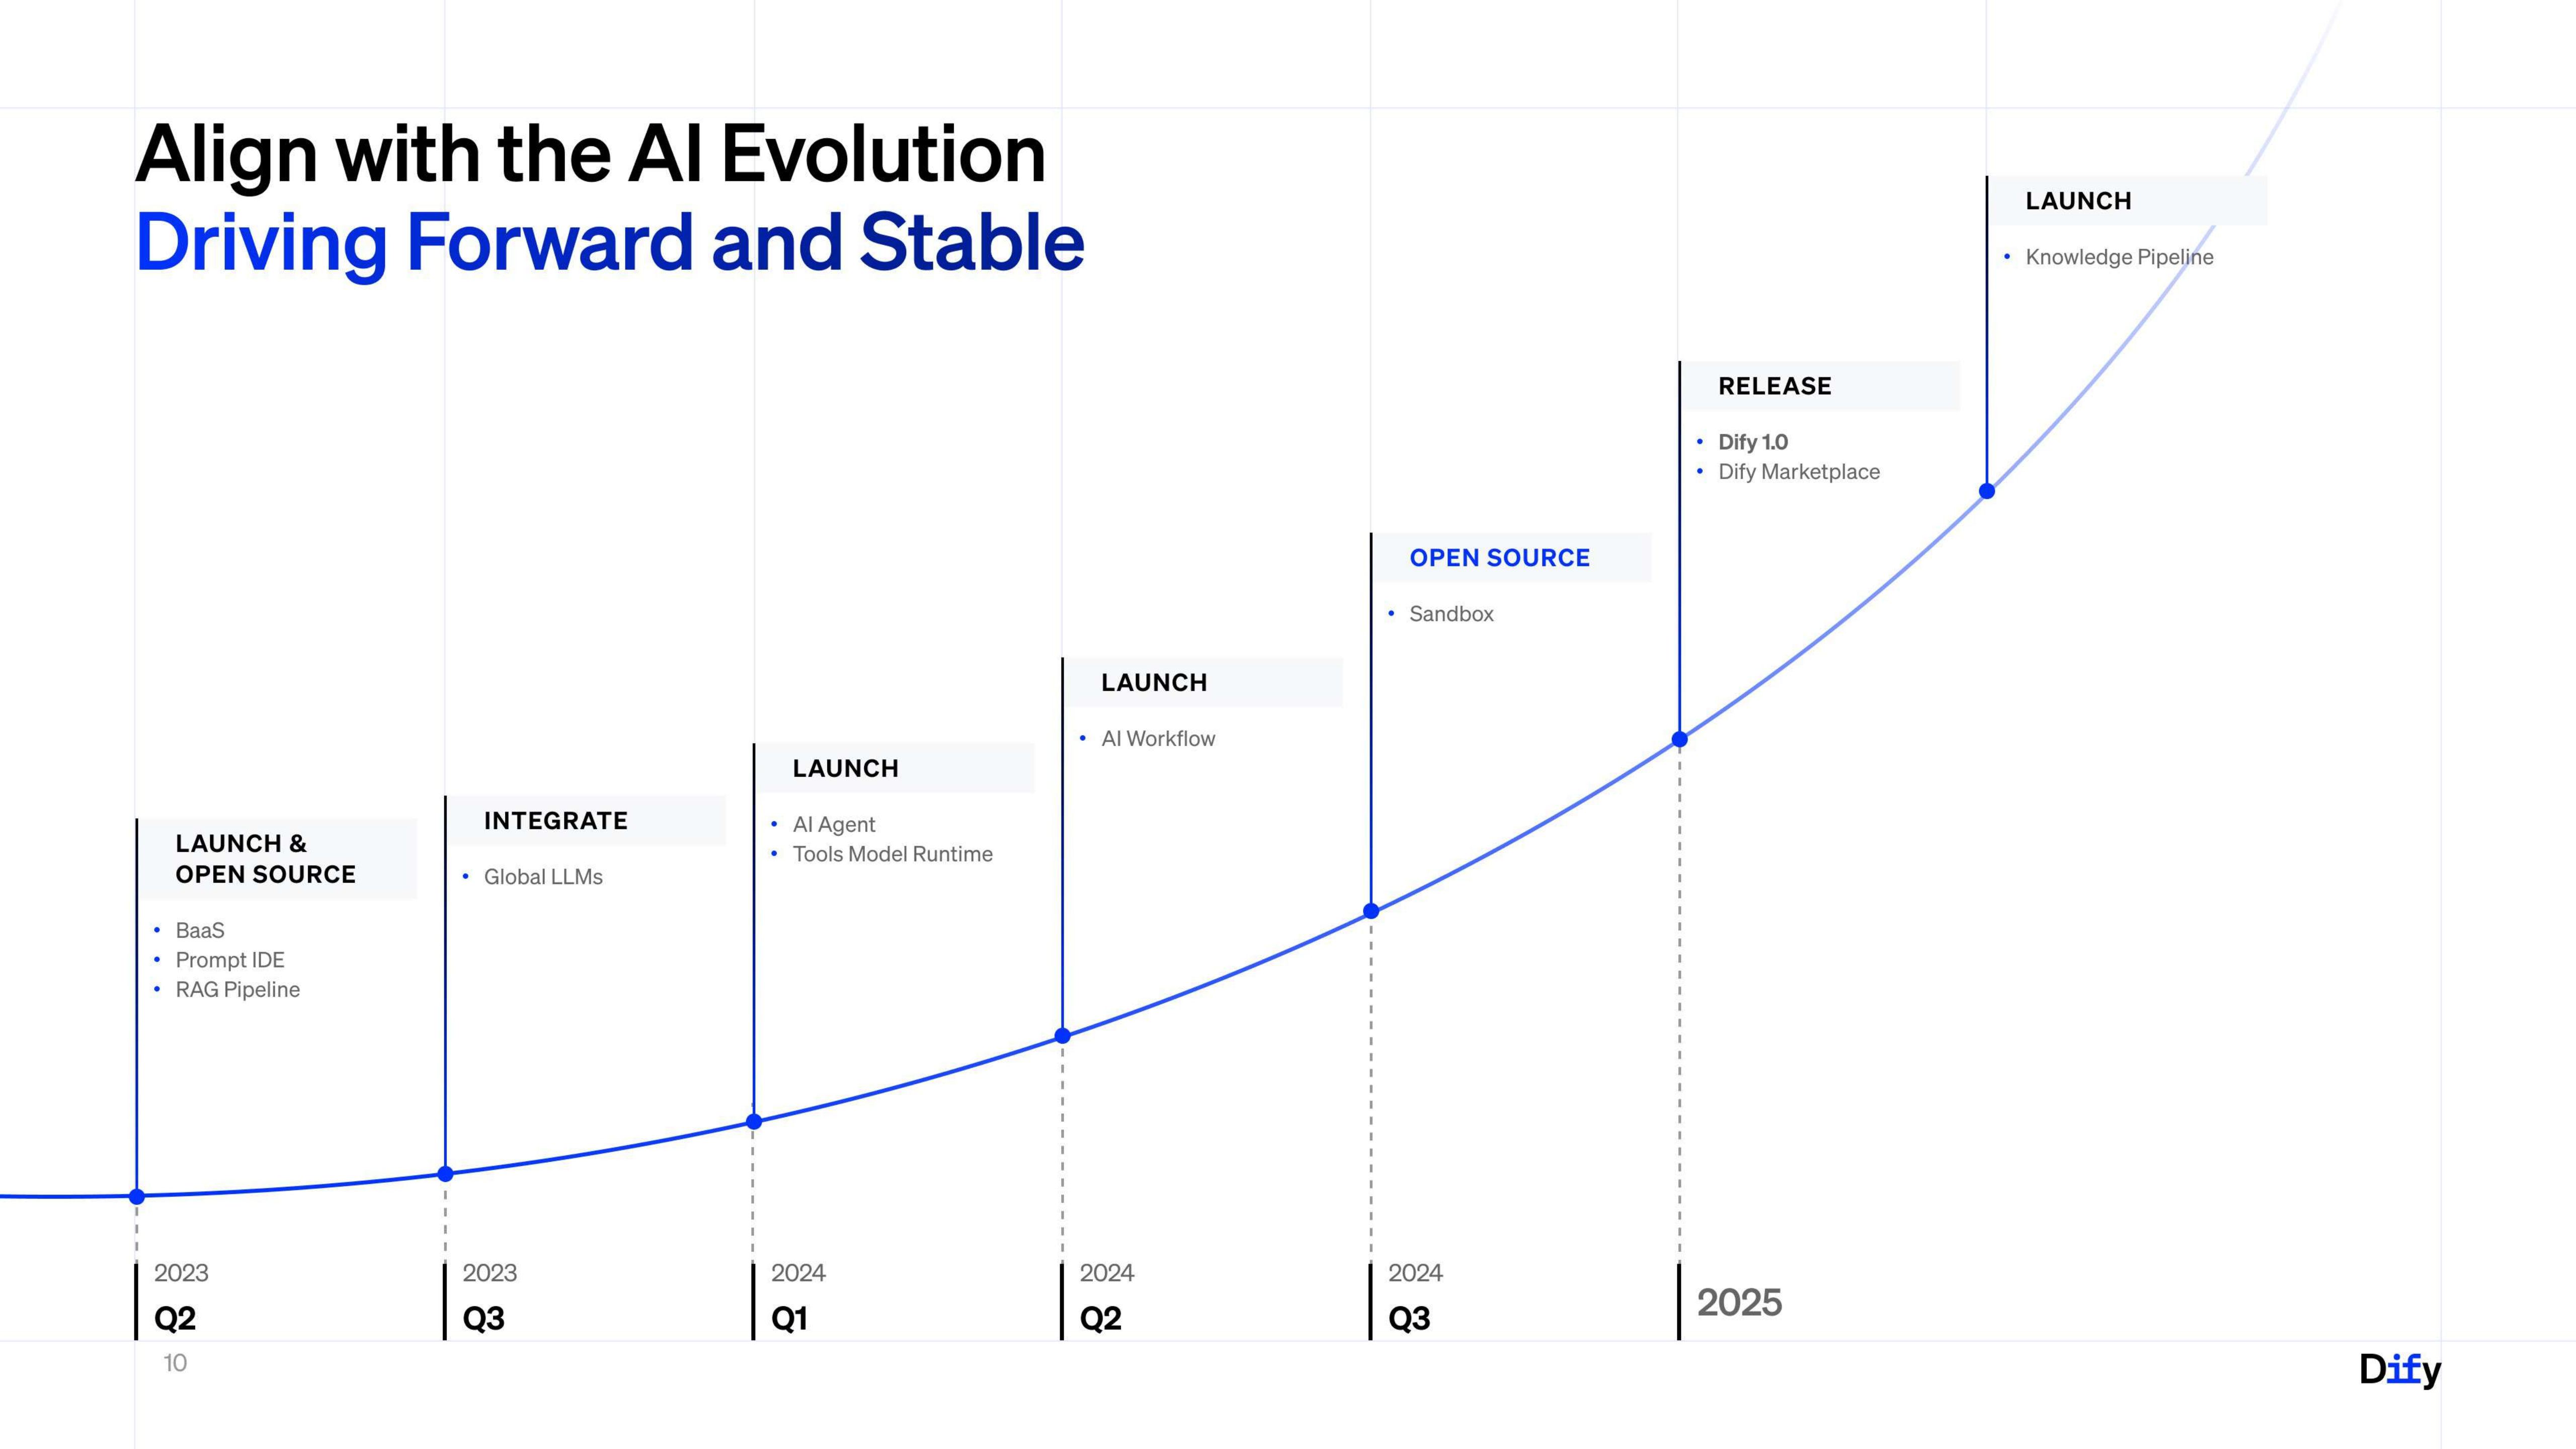
Service Partners

Our global partner network breaks down geographic and technical barriers to ensure your success with Dify, anywhere in the world.



Align with the AI Evolution

Driving Forward and Stable



2023
Q2

2023
Q3

2024
Q1

2024
Q2

2024
Q3

2025

Agentic Workflow

Drag and drop, orchestrate the nodes, and turn your AI vision into reality.



Cross-Multimodal Support

Optimal AI models (LLM, TTS, STT, etc.) for every task and business need.



Precise Memory Management

Conversation variables store temporary data, enabling multi-round dialogues.



Connect, Share, and Collaborate

Natively embed workflows within workflows. Share with the community and collaborate with your team.



Fault-Tolerant & High-Performance

Error handling, featuring backup branches and retries ensures workflow continuity.

The screenshot displays the Dify AI Studio interface. At the top, the navigation bar includes 'Dify', 'Solar Studio', 'TEAM', 'EXPLORE', 'STUDIO', 'CHATBOT PLAYGROUND', 'KNOWLEDGE', 'TOOLS', and 'PLUGINS'. The main workspace shows a workflow with four nodes: 'DOC EXTRACTOR', 'LLM', 'CODE', and another 'LLM'. The 'LLM' node is highlighted with a blue border. A 'Test Run#5' panel is open on the right, showing a 'TRACING' view of the workflow execution. The tracing shows the following steps:

- START: 0.02 s ✓
- DOC EXTRACTOR: 0.27 s ✓
- LLM: 258 tokens · 0.12 s ✓
- CODE: 1.53 s ⚠

The 'CODE' node has an orange warning box: "Node exception, will automatically execute the fail branch. The node output will return an error type and error message and pass them to downstream." Below this, the 'INPUT' and 'OUTPUT' sections are visible. The 'INPUT' section shows a JSON object:


```
{ "query": "What model are you?" }
```

. The 'OUTPUT' section is partially visible, showing a JSON object:

```
{
```

Knowledge Pipeline

Visually customize your data processing flow for better accuracy and relevance.



Connect Various Data Sources

Upload local files in various formats, import pages from Notion, sync content from websites, or connect to external knowledge base.



Get Your Data LLM-ready

Prepare your data for optimal LLM performance with Dify's RAG Pipeline—extract, process, transform and index seamlessly in various structures.



Modify, Organize, Update

Experience seamless knowledge base management with intuitive preview and editing, along with multiple APIs for effortless modification and updates.



Advanced Search Engine

Combine with LLMs for query classification, retrieve from multipath with metadata filtering and hybrid search, optimize results with Reranker, and cite sources for clear reference.


The screenshot displays the Dify Knowledge Pipeline interface. At the top, the navigation bar includes the Dify logo, user information (Solar Studio, TEAM), and menu items (EXPLORE, STUDIO, KNOWLEDGE, TOOLS, PLUGINS). The main workspace shows a workflow diagram with the following components:

- DATA SOURCE:** FILE UPLOAD (Upload local files) and NOTION (Sync Notion pages).
- DOC EXTRACTOR:** INPUT VARIABLE: File Upload / {x} documents File.
- GENERAL CHUNKER:** INPUT VARIABLE: Doc Extractor / {x} content File.
- KNOWLEDGE BASE:** Invoking large language models to answer questions or process natural language.

On the right, a configuration panel for the Knowledge Base is visible, showing settings for CHUNK STRUCTURE (General), INPUT VARIABLE (Doc Extractor / {x} content String), and INDEX METHOD (High Quality). The workflow is titled "Auto-Saved 21:02:35 · Unpublished" and includes a "Publish" button.

● MODEL AGNOSTIC

With Dify, you can unleash **any** large language models in just **a few clicks** ✨.



Connect with Global LLMs

Integrate any proprietary/open-source LLMs from dozens of inference providers and self-hosted solutions.



Dify Runtime

Integrate any model in just 10 minutes with a unified integration specification.



LLM API Load Balance

Optimize API requests distribution to increase model invocation capacity for fast response and stable performance.



Model Comparison

Compare performance from multiple LLMs to the same query for best selection based on your needs.

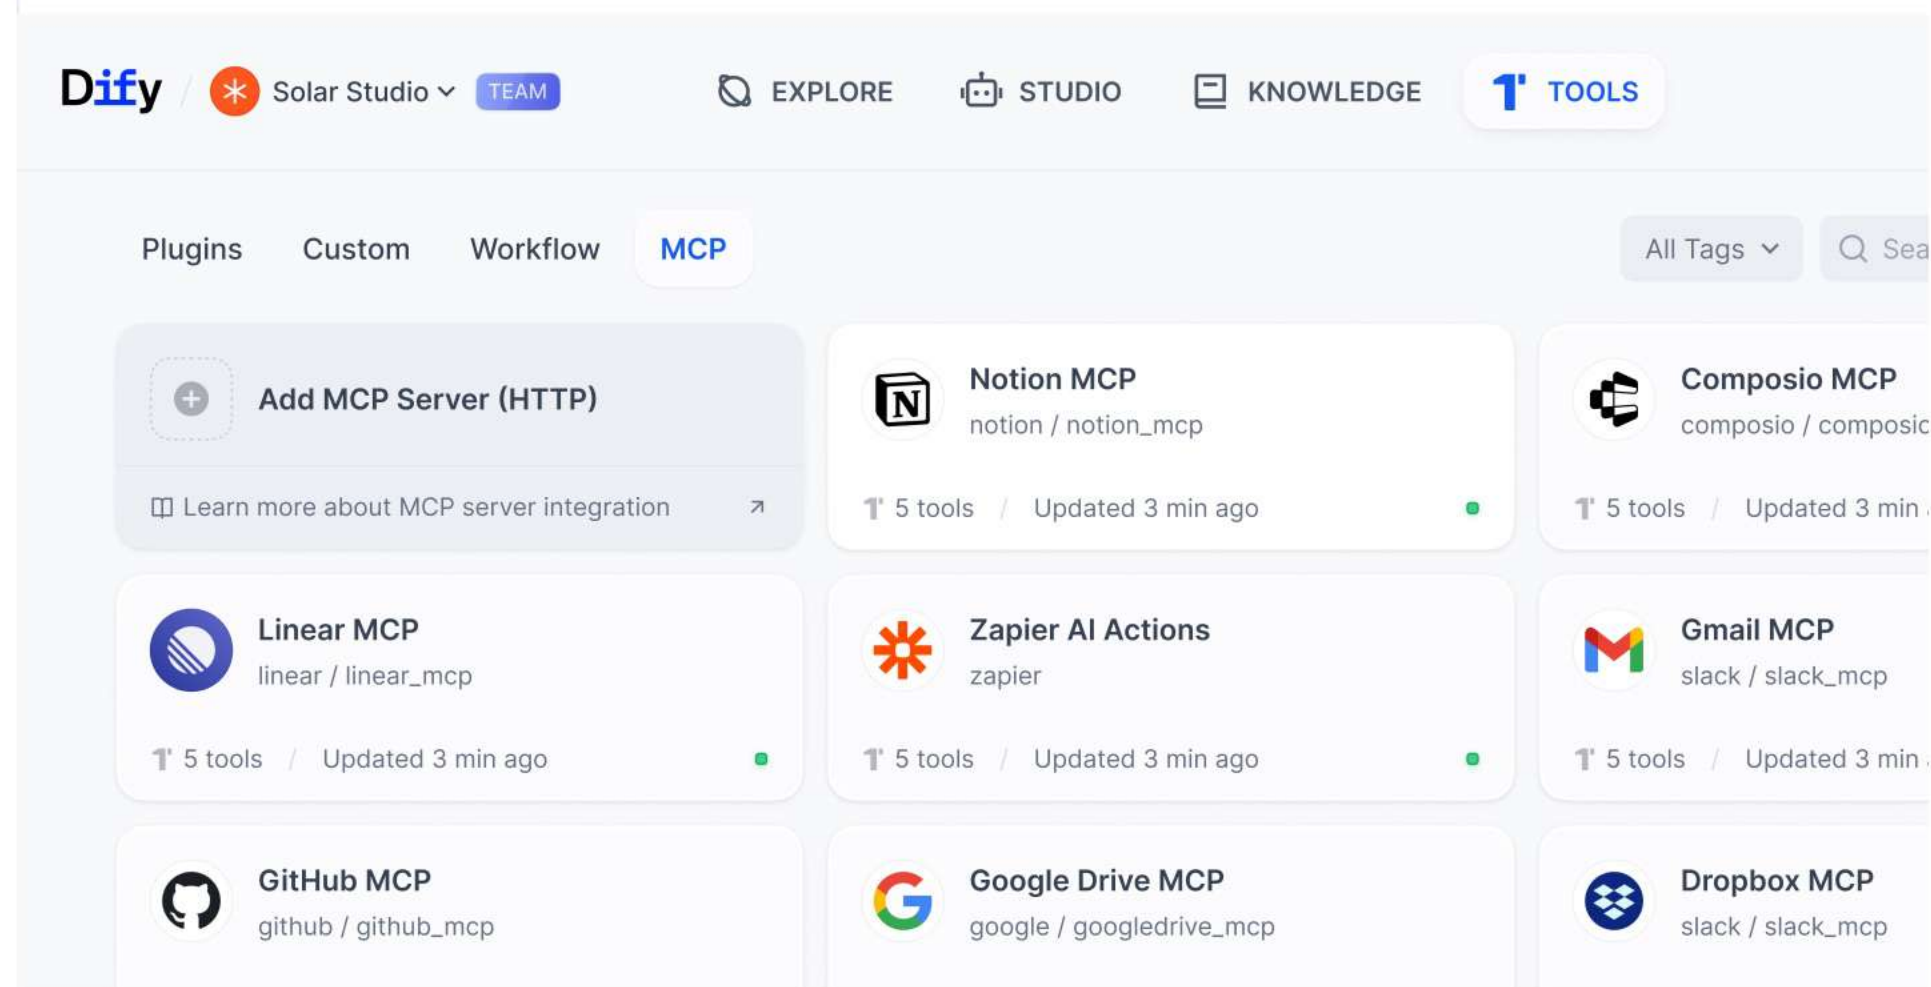
Built-in Two-way MCP Support

Seamlessly connect to external MCP servers directly or expose Dify-built agents or workflow as MCP servers.

Add an MCP Server

Connect to external MCP servers, such as Linear, Notion, Zapier to access their tools and resources. Unlock and extend Dify's capabilities by integrating third-party services, data sources and more.

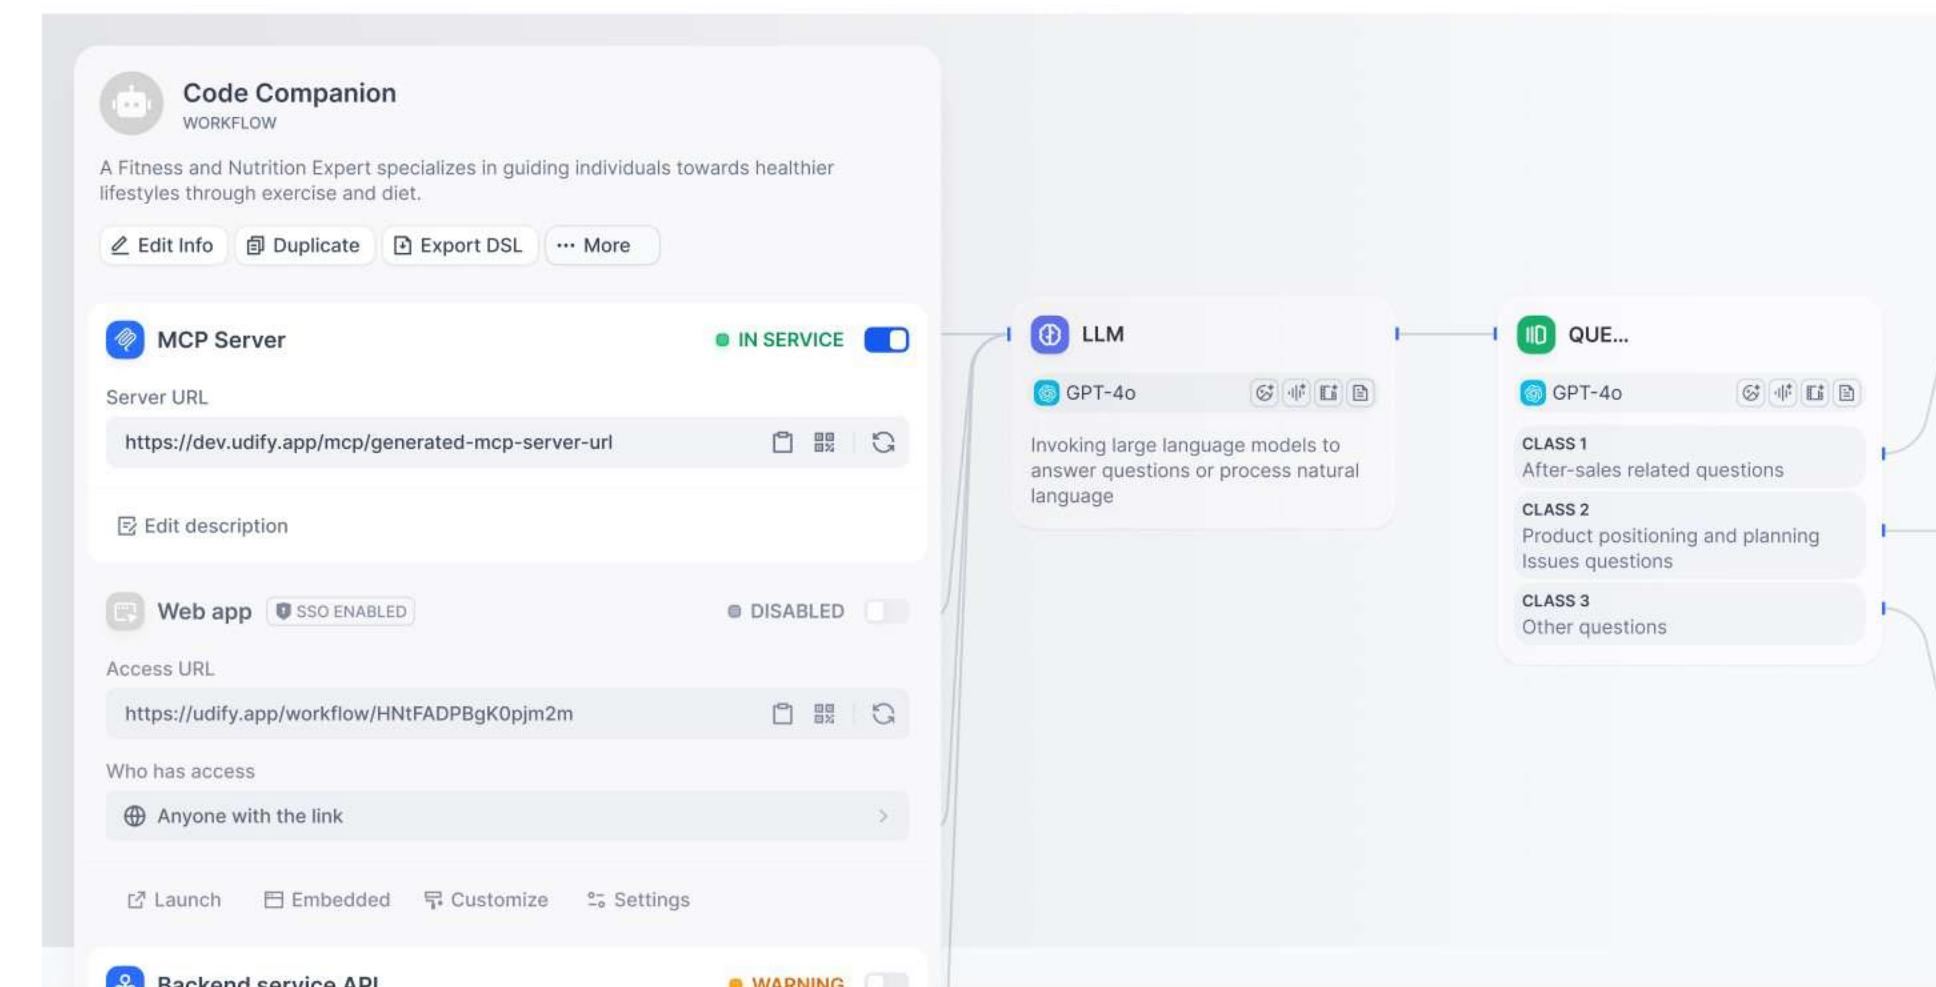
Example: Connect to Notion MCP server to let your agent read and write Notion pages



Publish as an MCP Server

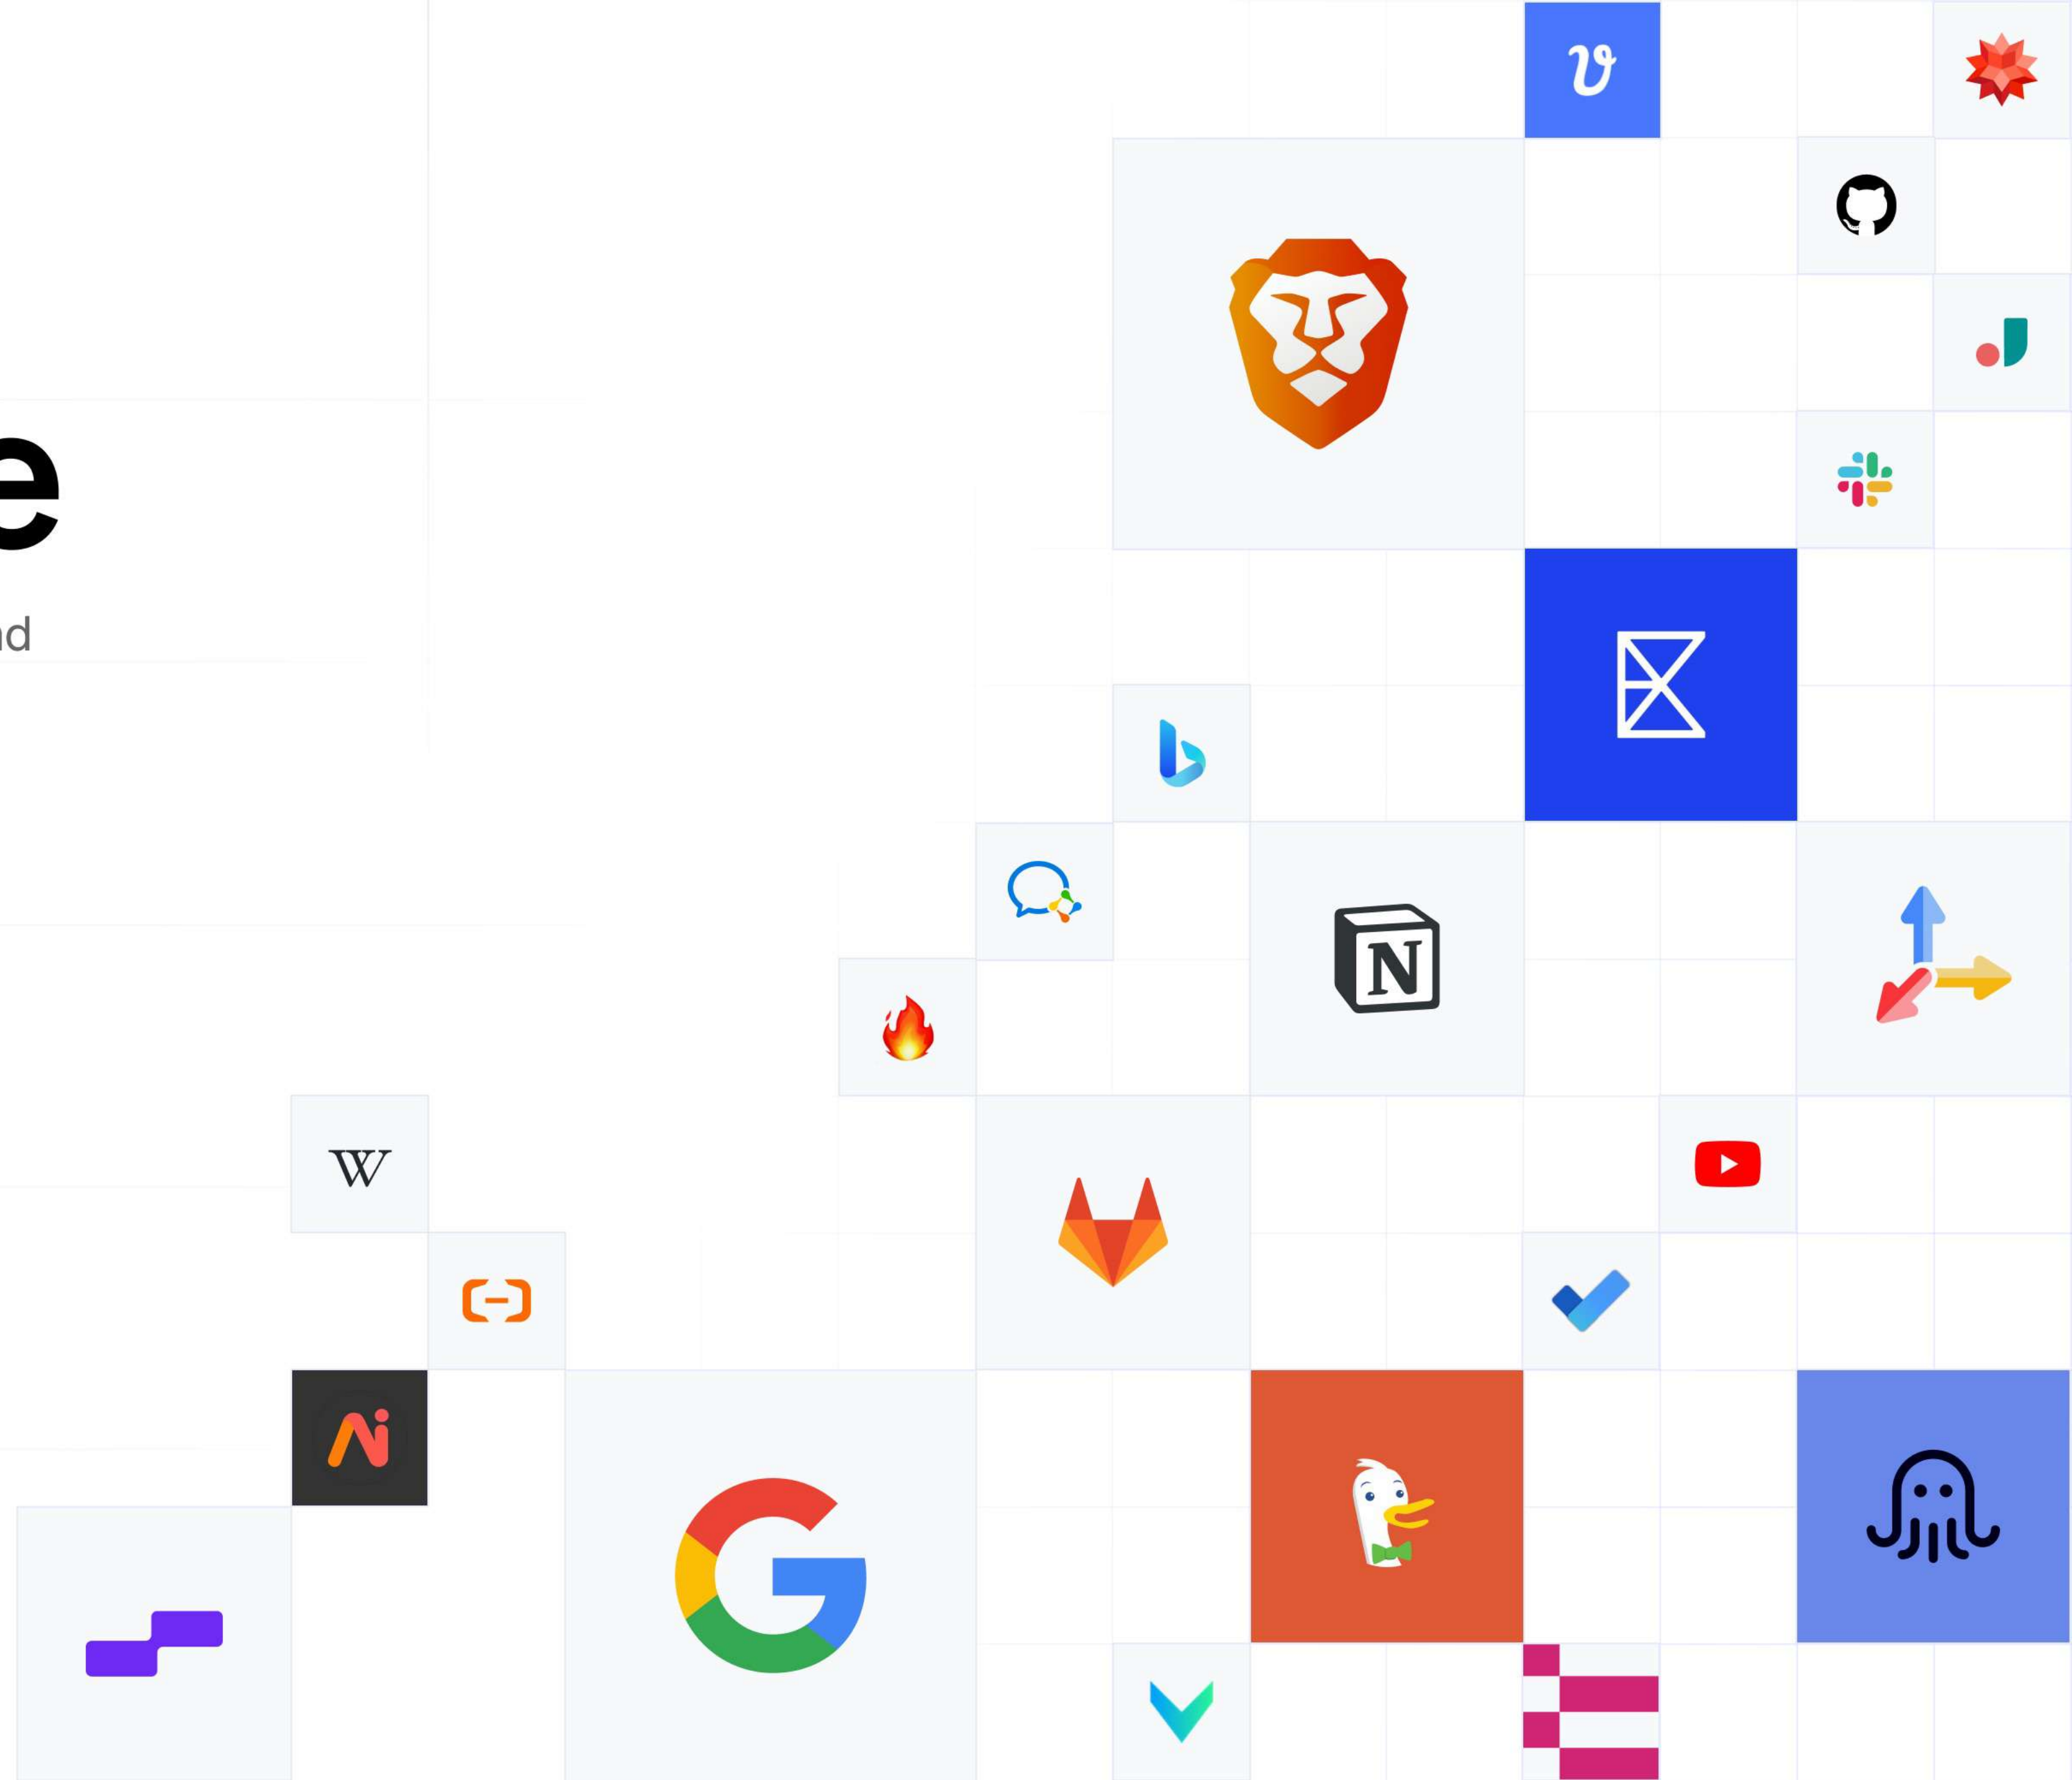
Expose Dify-built workflow or agent as a standard MCP endpoint. Turn it into a reusable MCP-standard server that tools such as Claude, Cursor, or any other MCP client can directly call.

Example: Publish a code review agent as an MCP server that Cursor can call directly



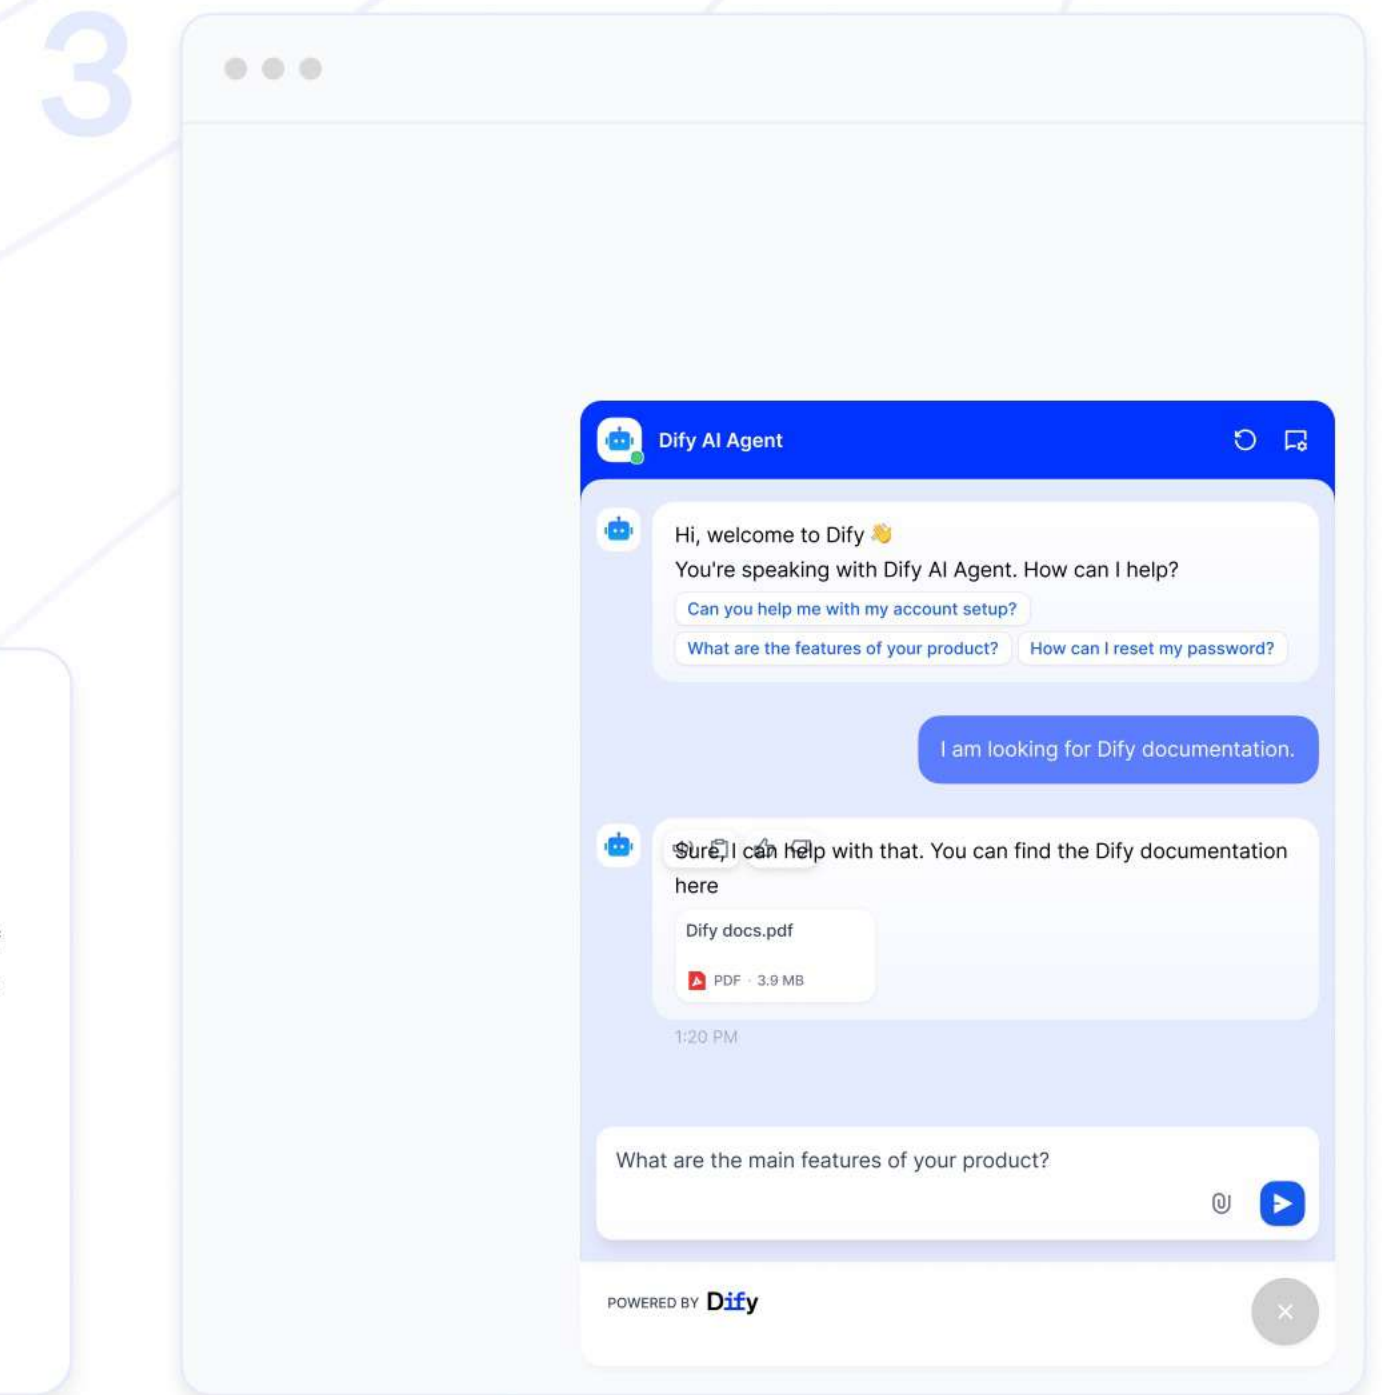
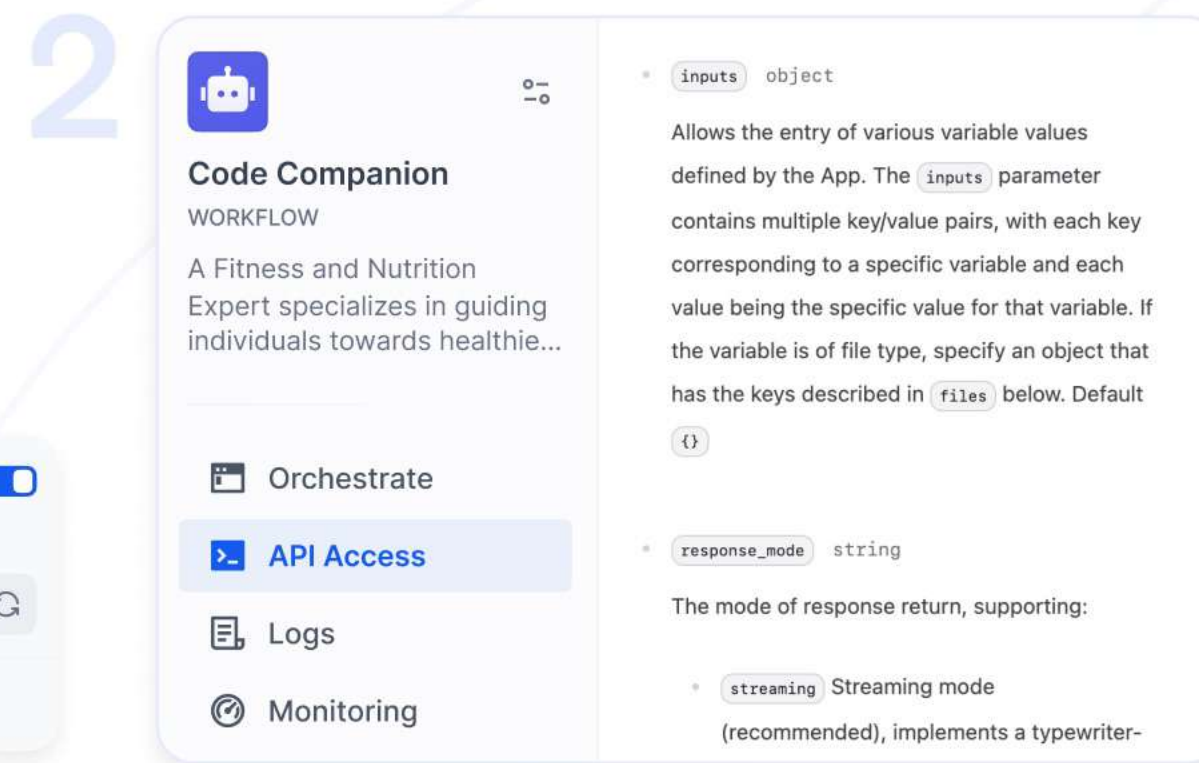
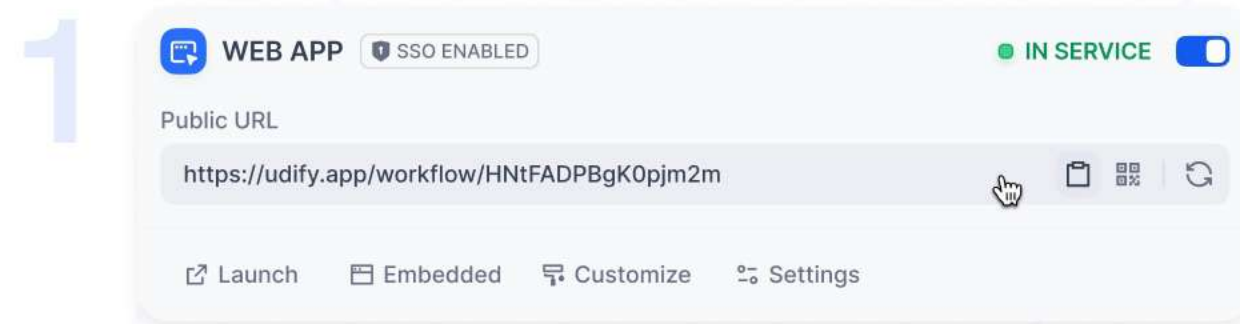
Marketplace

Dive into a growing ecosystem of powerful tools and integrations, built by Dify and global partners.



● DISTRIBUTION

Share it as **web app**¹
embed² it directly on your site
or **integrate**³ it into your system.



LLMOps & Monitoring

Gain complete AI application lifecycle tracking through Dify's native monitoring and seamless integrations with external LLMOps tools.

The screenshot displays the Dify Monitoring dashboard with several key components:

- Monitoring Header:** Includes a date selector for "Last 7 days".
- TOTAL MESSAGES:** A card showing a total of 250k messages.
- ACTIVE USERS:** A line chart showing 20k active users from January to June.
- TOKEN OUTPUT SPEED:** A line chart showing 1087 tokens per second from January to June.
- Tracing Panel:** A sidebar on the right with a toggle for "IN SERVICE" and a list of configured providers: LangSmith (IN USE), Langfuse, and opik.
- ANNOTATED RESPONSE Modal:** A pop-up window showing a user query and a bot response. The user query is "Let's play a decryption game today. You go first." The bot response is "Lara, the Caesar cipher is a simple substitution encryption technique, where each letter in the alphabet is shifted forward or backward a fixed number of positions. Please try shifting the letters back 13 positions." The modal includes "Save" and "Cancel" buttons, a "Remove this Cache" option, and a timestamp "CREATED AT 2023-03-21 10:00".

Month	Active Users
Jan	200
Feb	200
Mar	300
Apr	400
May	500
Jun	600

Month	Token/s
Jan	150
Feb	150
Mar	200
Apr	300
May	350
Jun	400

There's Always a Plan for Every Business Scope

Community

Open Source

- Core features released under the public repository
- Apache 2.0 License with additional conditions
- Community-driven support

SaaS


Start for Free Charged per Workspace/Month

- Everything out of the box
- Hosted infrastructure
- Up to 50 team members
- Flexible monthly billing, cancel at anytime

Enterprise

Annual Biling

- Private cloud / On-premise
- Enterprise-grade scalable deployment solutions
- Enterprise SSO & RBAC
- Maintenance and technical support



● FOR ENTERPRISE

Dify has made user data security, privacy protection, and compliance a top priority **since day one.**



Security & Access Control

SSO Integration

SAML, OIDC, OAuth2 support with enterprise systems (Okta, etc.)

Multi-Factor Authentication

TOTP-based two-step verification

Member Management

User/group management with SCIM 2.0 sync



Enterprise Management

Multi-Tenant Architecture

Isolated workspaces for data security

Centralized Console

Unified management hub

Enterprise APIs

Workspace and member management capabilities



Performance & Reliability

Load Balancing

Model API load distribution

High Availability

Kubernetes-based architecture

Easy Deployment

Helm Chart support



Branding Opportunity

Customization

Logo, name, and icon

Multi-Touchpoint

Consistent branding across all interfaces

Meet Some of Our Customers ★★★★★

Leading enterprises choose Dify to rapidly build and deploy production-ready AI application.

“ In this climate of perpetual beta, tools enabling rapid validation aren't just helpful, they're **existential**. For Volvo Cars, strategically navigating this AI frontier, this is where Dify delivers **indispensable value**.

V O L V O

GARY FENG
DIRECTOR OF AI • VOLVO CARS

“ What makes Dify stand out is its ability to democratize AI agent development. By combining powerful AI/ML capabilities on a no-code platform, its rapid deployment and intuitive interface make it highly accessible even for beginners, significantly accelerating citizen development.

RICOH

YOSHIAKI UMEZU
DIVISION GENERAL MANAGER • RICOH

Trusted by Industry Leaders

Real-world Results Powered by Dify

佰格.com

18K

hours saved / yr

- Company-wide chatbot saved 18,000 hours annually, equivalent to **10 full-time employees**
- Meeting minutes app developed in one month, saving an **additional 2,600 hours/year**

 MAERSK

19K+

employees served

- Dify-powered Q&A bot **serving 19,000+ employees across 20+ departments**
- Provides instant, unified information access breaking down data silos

ThermoFisher
SCIENTIFIC

300

hours saved / mo

- Saves approximately 300 hours monthly with **support ticket auto-filling workflow**
- **Single IT engineer** prototyped and deployed workflow handling hundreds of daily inquiries

From Business Challenge to Tangible Results

Case Study on Kakaku.com



Kakaku.com, Inc. (est. 1997) is a Japanese digital leader behind Kakaku.com (price comparison), Carview! (automotive), and Tabelog (restaurant reviews).

BEFORE • CHALLENGE

Limited AI Expertise

Drive innovation through enterprise-wide AI adoption but with limited AI specialists

Slow Development Cycles

Long development cycles despite promising POCs

Inaccessible AI Benefits

AI capabilities remained siloed in tech teams, preventing 95% of employees from leveraging AI in their daily work.

AFTER • SOLUTION

Self-Service AI

Empowered non-technical teams to build and deploy AI chatbots without writing code or waiting for engineering resources.

DEMOCRATIZED ACCESS

Instant Content Generation

Automated article titles and content creation with pre-built workflows, reducing time-to-market from months to hours.

PRODUCTION-READY IN HOURS

Enterprise-Wide Deployment

Launched company-wide financial assistant chatbot serving all employees with minimal IT involvement and zero custom development.

3-HOUR SETUP

Before



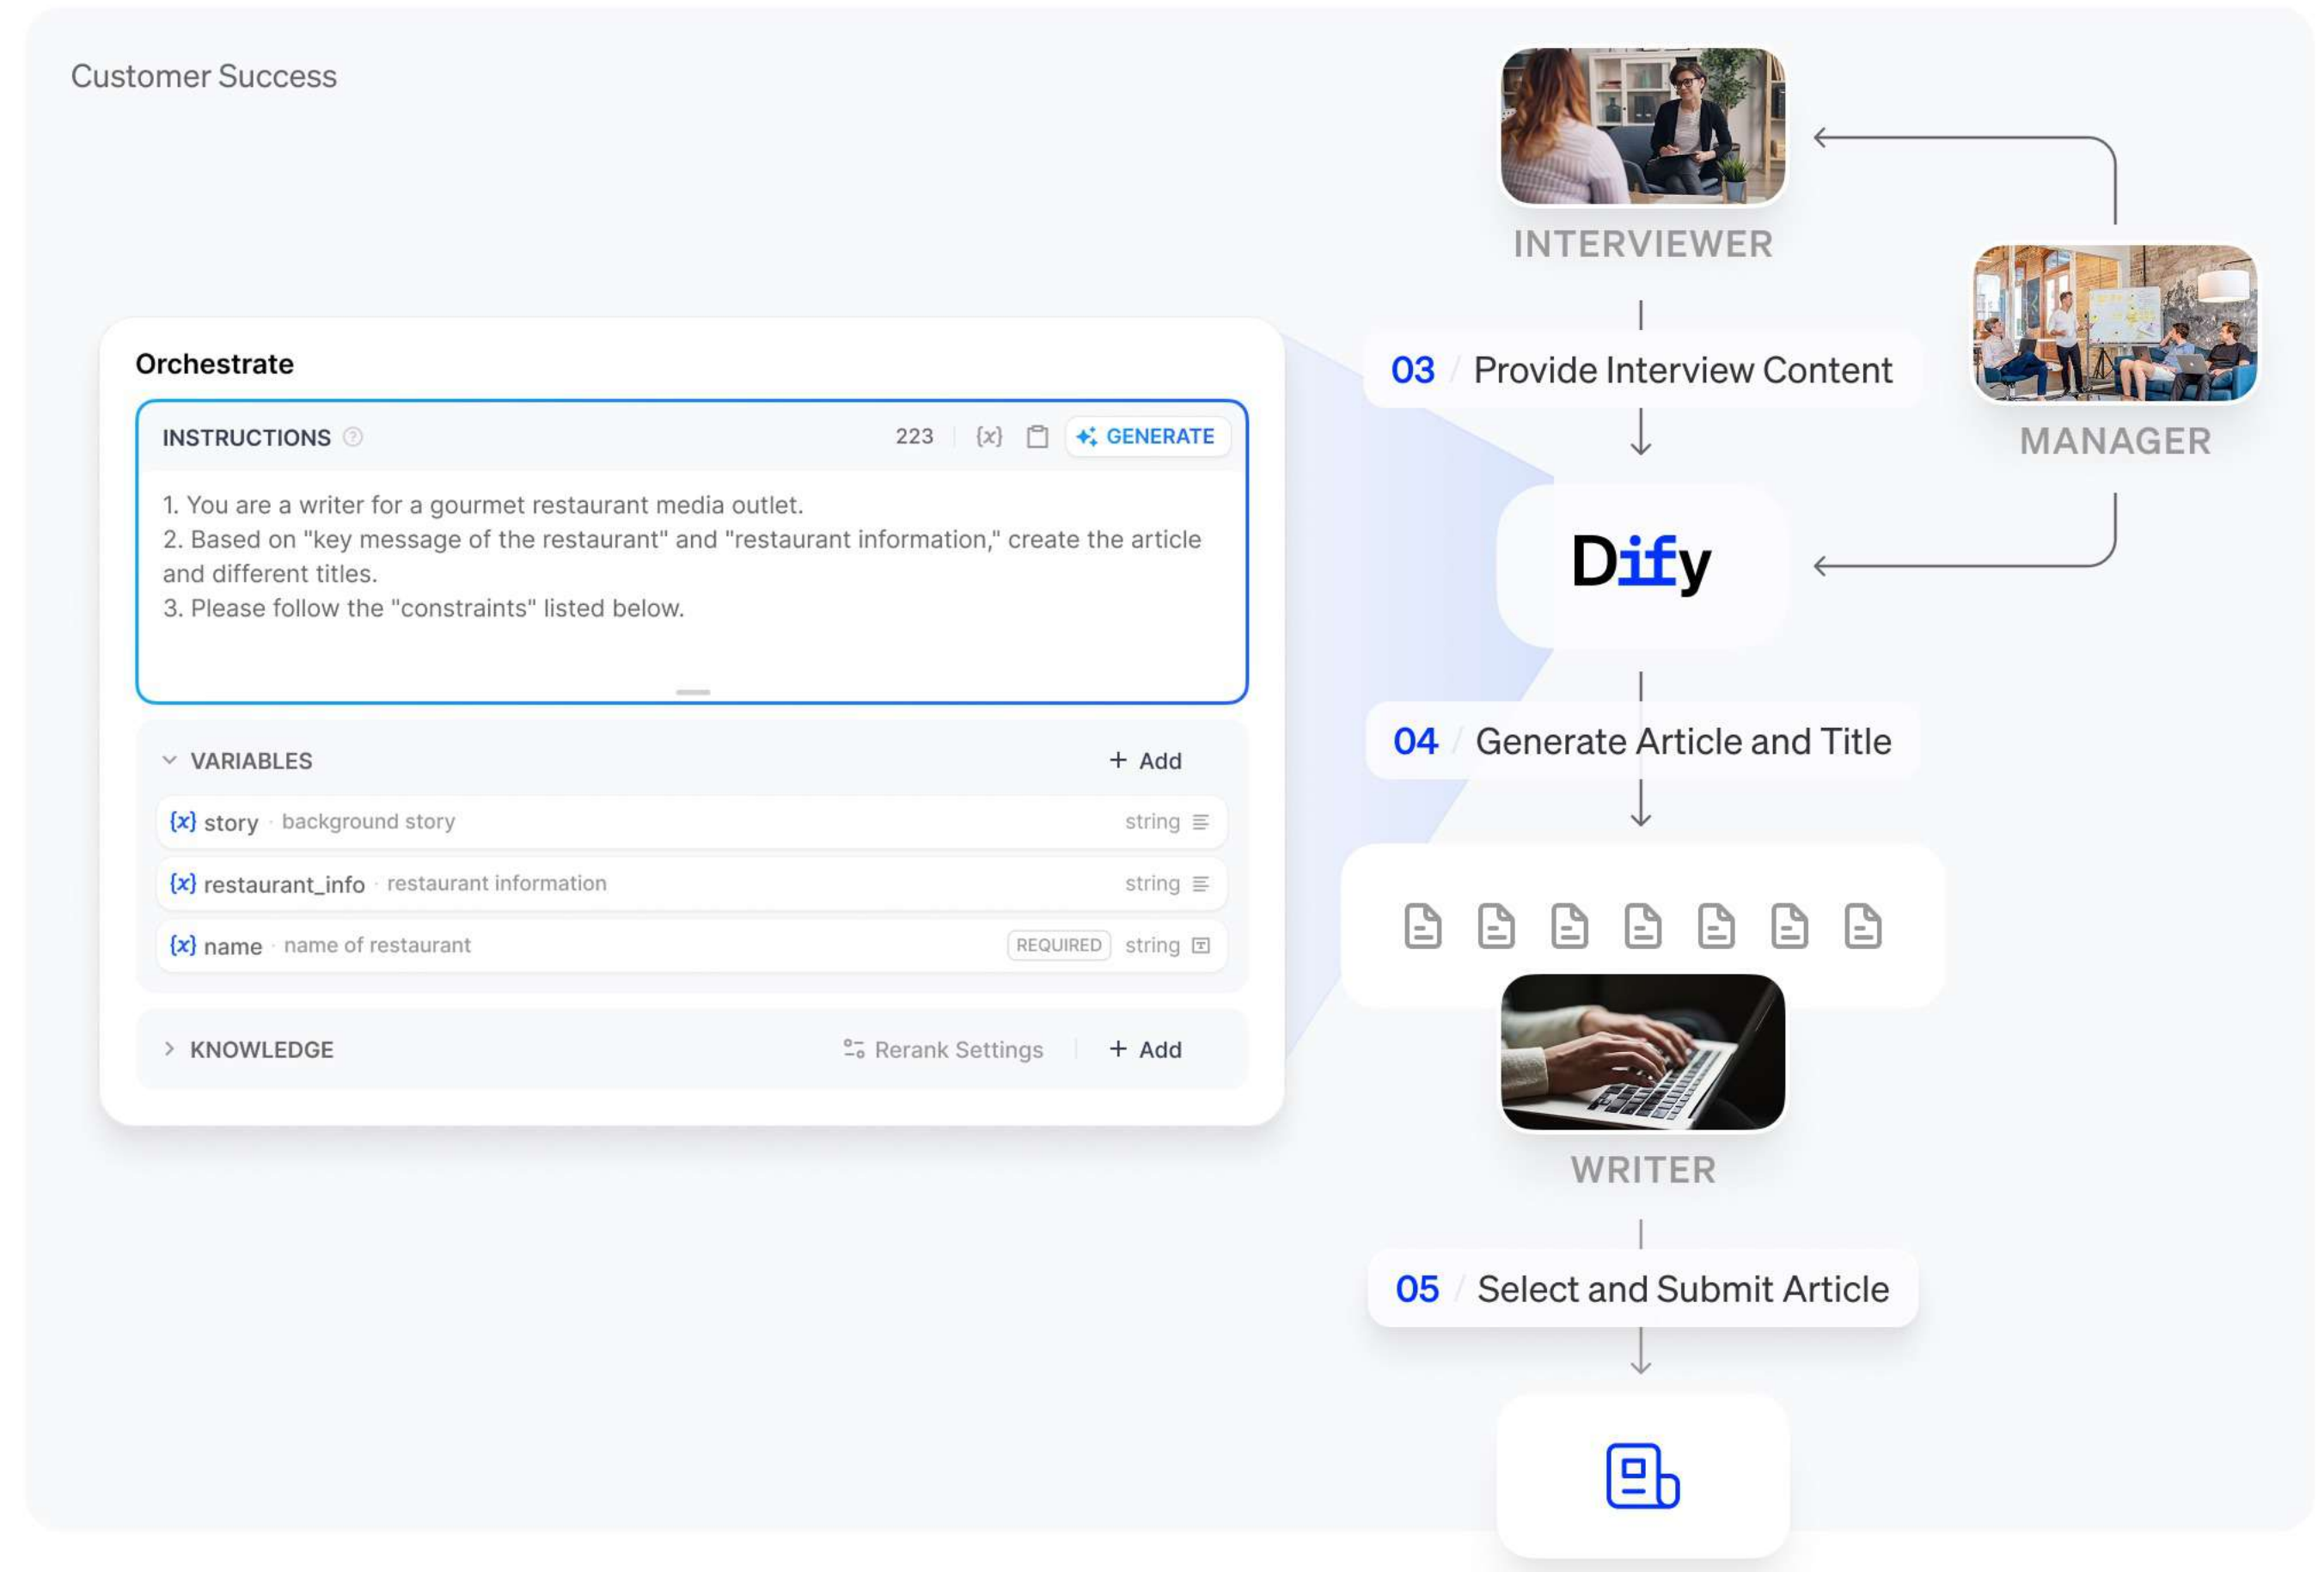
RESTAURANT



After



RESTAURANT



From Business Challenge to Tangible Results

Case Study on Maersk Digital Transformation



Maersk is a global leader in integrated logistics, managing 20% of world's container shipping across 130+ countries.

BEFORE • CHALLENGE

Critical Information Bottleneck

Employees struggled to access scattered information, causing delays per audit cycle and operational inefficiencies.

Manual Process Overload

Contract reviews and compliance audits required extensive manual effort, creating documentation gaps and risks.

AFTER • SOLUTION

Enterprise AI Knowledge Hub

DEMOCRATIZED ACCESS

Deployed an intelligent chatbot serving 19,000+ employees across all departments, providing instant access to company knowledge, policies, and historical data with natural language queries.

AI-Powered Audit Automation

PRODUCTION-READY

Implemented automated contract review system with AI knowledge base analyzing past contracts, streamlining compliance processes and eliminating manual documentation bottlenecks.

RESULTS

90%

Reduction in audit cycle time from 40+ hours to 4 hours

19K+

Employees with instant access to enterprise knowledge

\$2.4M

Annual savings in operational efficiency and productivity gains

100%

Compliance documentation accuracy with automated reviews

Before

THE CHALLENGE

Manual & Overwhelming

Maersk team needs to process and manage thousands of contracts every day, which brings them management pressure and challenges.

- ⊗ Slow manual processing
- ⊗ High management pressure
- ⊗ Operational challenges

After

THE TRANSFORMATION

Automated & Efficient

 **Boost Contract Processing Efficiency**
Automated workflows drastically cut processing time and accelerate business response.

 **Boost Contract Processing Efficiency**
Automated workflows drastically cut processing time and accelerate business response.

Dify

WEBSITE
dify.ai

GITHUB
github.com/langgenius/dify

X
x.com/dify_ai

DISCORD
discord.gg/FngNHpbcY7

LINKEDIN
linkedin.com/company/langgenius

EMAIL
business@dify.ai