

GenAI - Empowering the Developer Journey at Every Step

Developer Productivity Engineering Summit
9/24/24



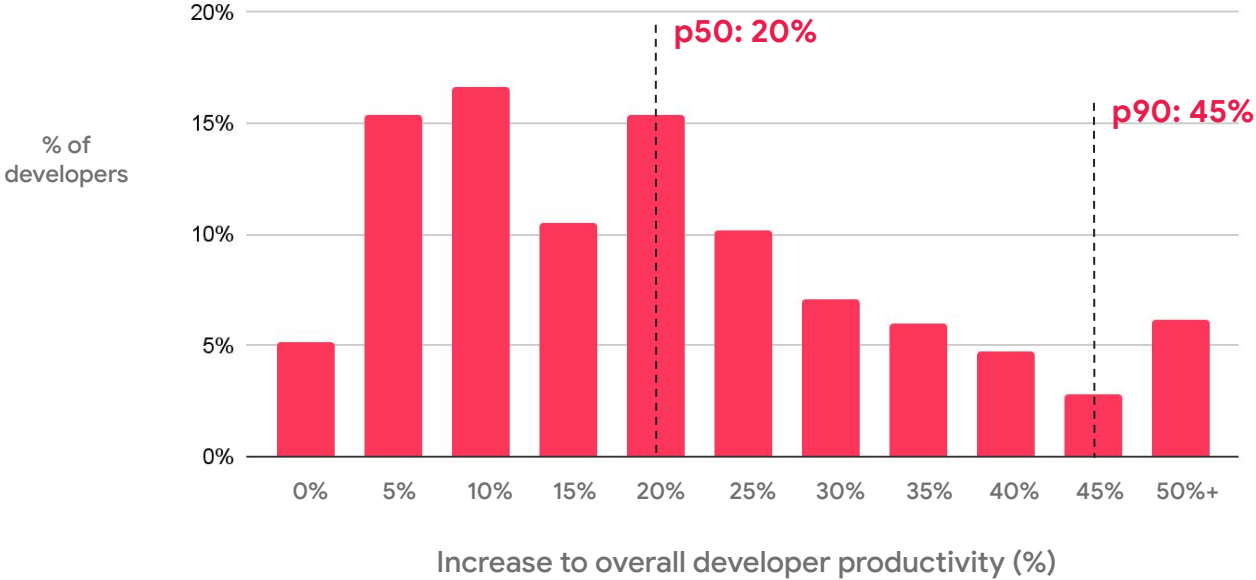
Byron Grogan

Software Engineer, Developer Platform
Airbnb

Please list any improvements or product launches (in the last 6 months) that have positively impacted your developer workflow

Airbnb's DevX Survey

Developers say AI tools have increased their overall productivity by 20% on average



96% of developers said it had a >0% positive impact*

*Response options ranged from 0 - 50%



Chat



Copilot



Platform & Integrations

Chat

Chat Use Cases

Coding

Writing Tests

Debugging

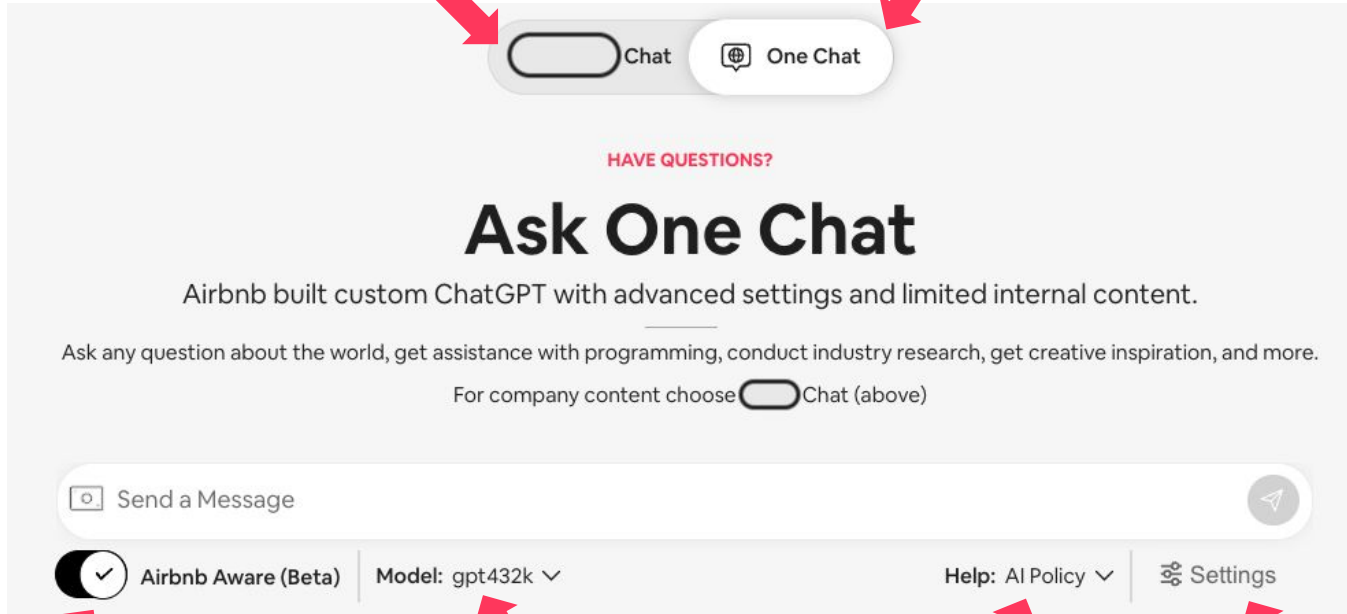
Refactoring

Writing Docs

Engineering Support

Use vendor chat for 100% Airbnb specific questions: workflows, policy documents, meeting notes, and more.

Use One Chat for everything else, including code!



Internal knowledge base aware (retrieval augmented generation - RAG)

Pick your model

Airbnb's AI policy

Tune settings: system message, "temperature", token limit

Example One Chat conversation

Example chat



Chat history



+ New chat

Favorites ^

- ☆ "MySQL Configuration Upg

Today ^

- "Understanding Artifact Up

Previous 7 days ^

- "Viewing Helm Release File
- "Page Cache and Cgroups"

6:26:29 PM
how to implement retry in bazel test in code



6:26:29 PM

To implement retry logic in Bazel tests, you can use the `--flaky_test_attempts` flag. Then Bazel will rerun it up to three times if it fails.

Here is an example of a flaky test:



6:32:12 PM
what is shard_count in go_test



6:32:12 PM

The `shard_count` in the context of a `go_test` is the number of partitions that a set of tests will be run on. This is useful when you want to parallelize test execution to reduce the overall test runtime. Each shard can be run in parallel, ideally on a separate processing unit, to speed up the overall test execution.

6:36:21 PM
what is embed in go_test



6:36:21 PM

The `embed` attribute in the `go_test` rule of Bazel is used to include

Airbnb's One Chat

User prompt



6:36:21 PM
what is embed in go_test



6:36:21 PM

The `embed` attribute in the `go_test` rule of Bazel is used to include

Default system message



maxTokens

3000

The maximum number of tokens to generate in the chat completion. The total length of input tokens and generated tokens is limited by the model's context length.

Temperature

0.8

What sampling temperature to use, between 0 and 2. Higher values like 0.8 will make the output more random, while lower values like 0.2 will make it more focused and deterministic.

stop

String where the API will stop generating further tokens.

System Message

You are an AI assistant for Airbnb employees. Help make them productive with details, helpful information, and tips.

****bold**** *_italics_* ~~~~strike~~~~ [link](https://)

Chat Feature Comparison

	One Chat Web Interface	Copilot Chat	One Chat IDE Plugins
Integrated into the IDE	✗	✓	✓
Pulls context from active code	✗	✓	✓
Airbnb-aware RAG	✓	✗	✓
Choose your model	✓	✗	✓

Retrieval Augmented Generation (RAG)

How Airbnb's One Chat is "Airbnb aware"

Enter RAG (Retrieval Augmented Generation)

Tell the model what it needs to know

- Find relevant context:
 - Articles, code, etc. that are related
- Combine the context into the LLM prompt
- “Relevance” can be achieved in many ways:
 - If you already have search + preview system, reuse it
 - Embedding-based retrieval: use LLMs to find “relevant” snippets
 - Whatever happened recently in the system
 - Code / docs nearby in the filesystem

What is Onebrain?

Before RAG

Onebrain is an AI-driven software solution developed by Airbnb that aims to improve efficiency, productivity, and decision-making across the company. It integrates data from internal sources and uses machine learning algorithms to analyze and make sense of the data. This helps Airbnb employees make better decisions and streamline processes by providing insights, predictions, and recommendations based on the analyzed data.

After RAG

You are an assistant for Airbnb employees. You know a series of facts about Airbnb:


{output of internal search}

What is Onebrain?


OneBrain is a Python package and command line interface (CLI) developed by Airbnb. It is a versatile tool designed to power a scientific project ecosystem, providing integration with Redspot and other functionalities. OneBrain serves as a comprehensive platform for managing and executing various tasks within the Airbnb development environment. It helps Airbnb employees streamline their workflows, automate processes, and collaborate more effectively on projects.

Generate iOS project

Without Airbnb Aware

 12:22:42 PM 📄 🔗 ✎

How do I generate an iOS project?

 12:22:42 PM 👍 💬

Creating an iOS project involves using Xcode, which is the official integrated development environment (IDE) for iOS development. Here are the general steps:

- 1. Download Xcode:** You can download Xcode for free from the Mac App Store. You will need a Mac since Xcode is not available on other operating systems.
- 2. Start a New Project:** Open Xcode and go to "File" > "New" > "Project".
- 3. Select a Template:** A new window will pop up with several templates. If you're creating a standard iOS app, you'll probably want to choose "Single View App".
- 4. Configure Your Project:** Next, you'll need to fill out some details for your project:
 - **Product Name:** This is the name of your app.
 - **Team:** If you're part of an organization in Apple's Developer Program, you would select it here. If not, you can select "Add Account..." to sign in with your Apple ID.
 - **Organization Name:** This can be your personal name or your company's name.
 - **Organization Identifier:** This is often a reverse domain (like "com.yourCompanyName").
 - **Bundle Identifier:** This is automatically generated based on your product name and organization identifier.
 - **Language:** Choose Swift or Objective-C. If you're new to iOS development, you probably want to choose Swift.
 - **User Interface:** Choose Storyboard or SwiftUI. If you're new to iOS development, SwiftUI might be easier to start with.
- 5. Save Your Project:** Choose where you want to save your project on your computer, then click "Create".
- 6. Explore the App Structure:** A new project in Xcode comes with some default files that you can explore to understand the structure of an iOS project.

Remember, developing an iOS app requires a good understanding of either Swift or Objective-C, as well as knowledge of Apple's development frameworks. If you're new to programming or iOS development, you might want to start with some tutorials or courses.

RAG switch




📄 🔗 ✎


✕ Airbnb Aware (Beta) Model: gpt432k ^ Help: AI Policy ⚙️ Settings

Generate iOS project

With Airbnb Aware

 11:54:00 AM 📄 🔗 ✎

How do I generate an iOS project?

 11:54:00 AM 👍 💬

To generate an iOS project at Airbnb, you can use Crane, which is a tool used to generate an Xcode project. The `-m` argument accepts a whitespace-separated list of modules surrounded by quotes. For example, you can use `bear crane -m "SharingFeature SharedPDPlugin"` to generate a focused project with both SharingFeature and SharedPDPlugin loaded in the Xcode project. For more details, you can refer to the following articles:

- [How do I generate an Xcode project with code for both a feature and a plugin?](#)
- [History of iOS project generation](#)
- [What's Crane and how do I use it?](#)

Keep in mind that after generating the project, you might need to regenerate it when you switch branches or when there are changes in the source files. You can automate this process using a post-checkout script included in the repository, as explained in [Do I have to regenerate Buck-generated Xcode projects every time I switch branches?](#)

RAG switch



📄 📧

Airbnb Aware (Beta) | Model: gpt432k ⌵ | Help: AI Policy ⌵ | ⚙️ Settings

Copilot

IDE Autocomplete

Copilot Stats, according to GitHub

Improving developer satisfaction & Conserving mental energy

74% Focus on more satisfying work

88% Feel more productive

96% Faster with more repetitive tasks

Development speed

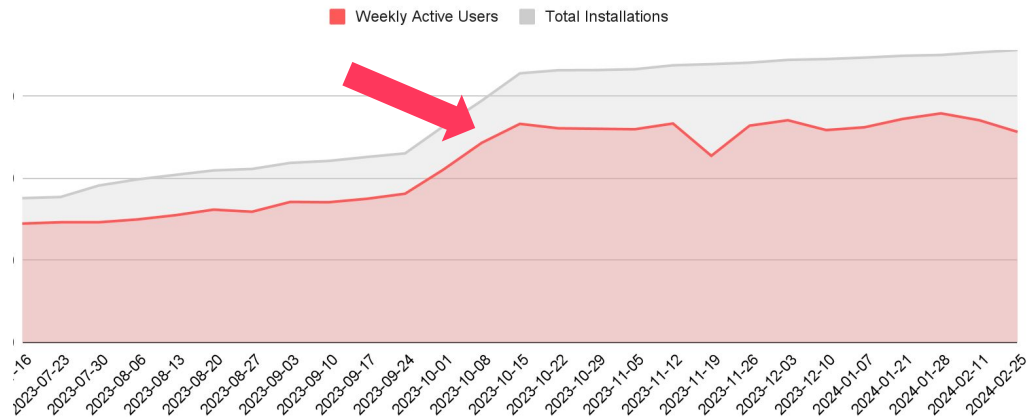
95 Developers given the same task

55% Less time to finish

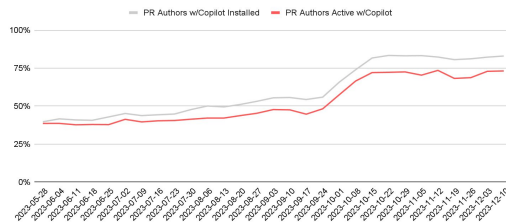
Copilot

- 80% of weekly PR authors have Copilot installed. 70% actively using*
- 75% of PRs are from authors with Copilot installed

Total Installations & Weekly Active Users



% of PR Authors with Copilot



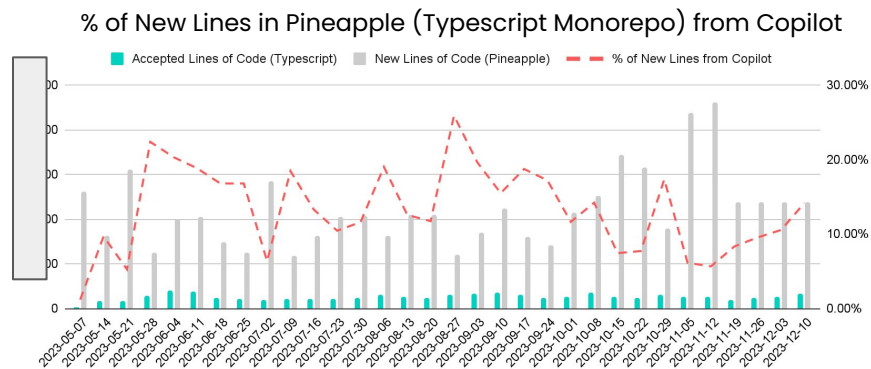
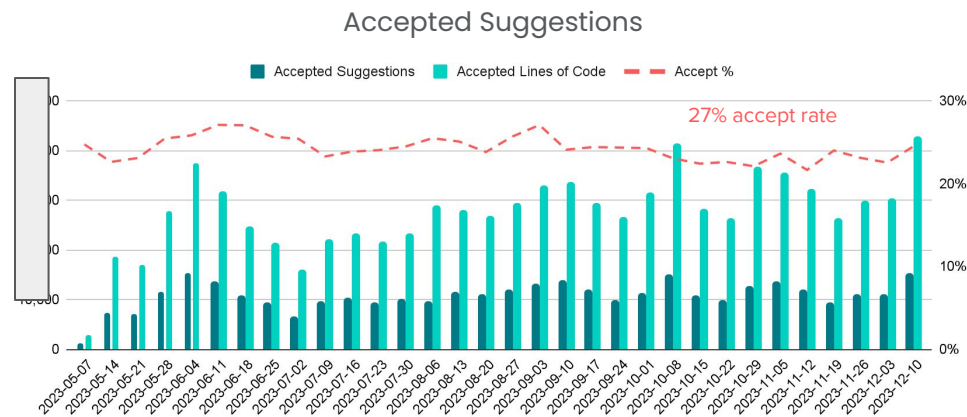
Total PRs from Authors with Copilot



* Active users as reported by GitHub

Copilot Engagement VS Code Only

- 13k weekly accepted suggestions
- 27% accept rate
- 16% of new lines of code for Web from Copilot*

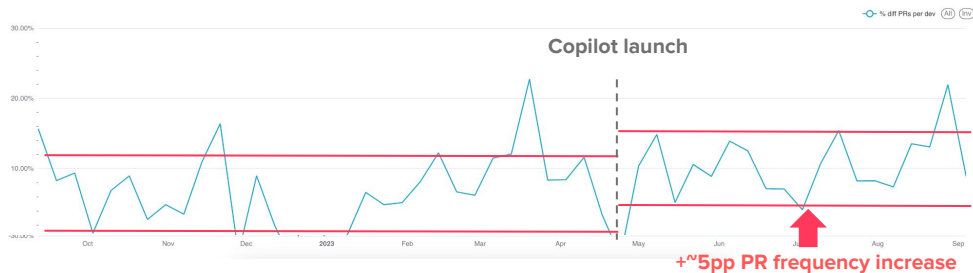


* % of new lines from Copilot estimated by accepted lines of typescript / new lines of code committed to our Typescript monorepo. Since VS Code is used for more than just that repo this is likely overcounting.

Impact on Productivity

- Copilot developers submit PRs at ~7-15% higher frequency than non-Copilot developers (up ~5pp since Copilot launch)

PR frequency of developers with Copilot vs without (% difference)



According to outside research

- Copilot users [completed a task 56% faster](#) than those without. 35% self-estimated increase in productivity.
- [97% of developers](#) said Copilot made them more efficient and productive.

Platform & Integrations



Innovate beyond Chat and Copilot

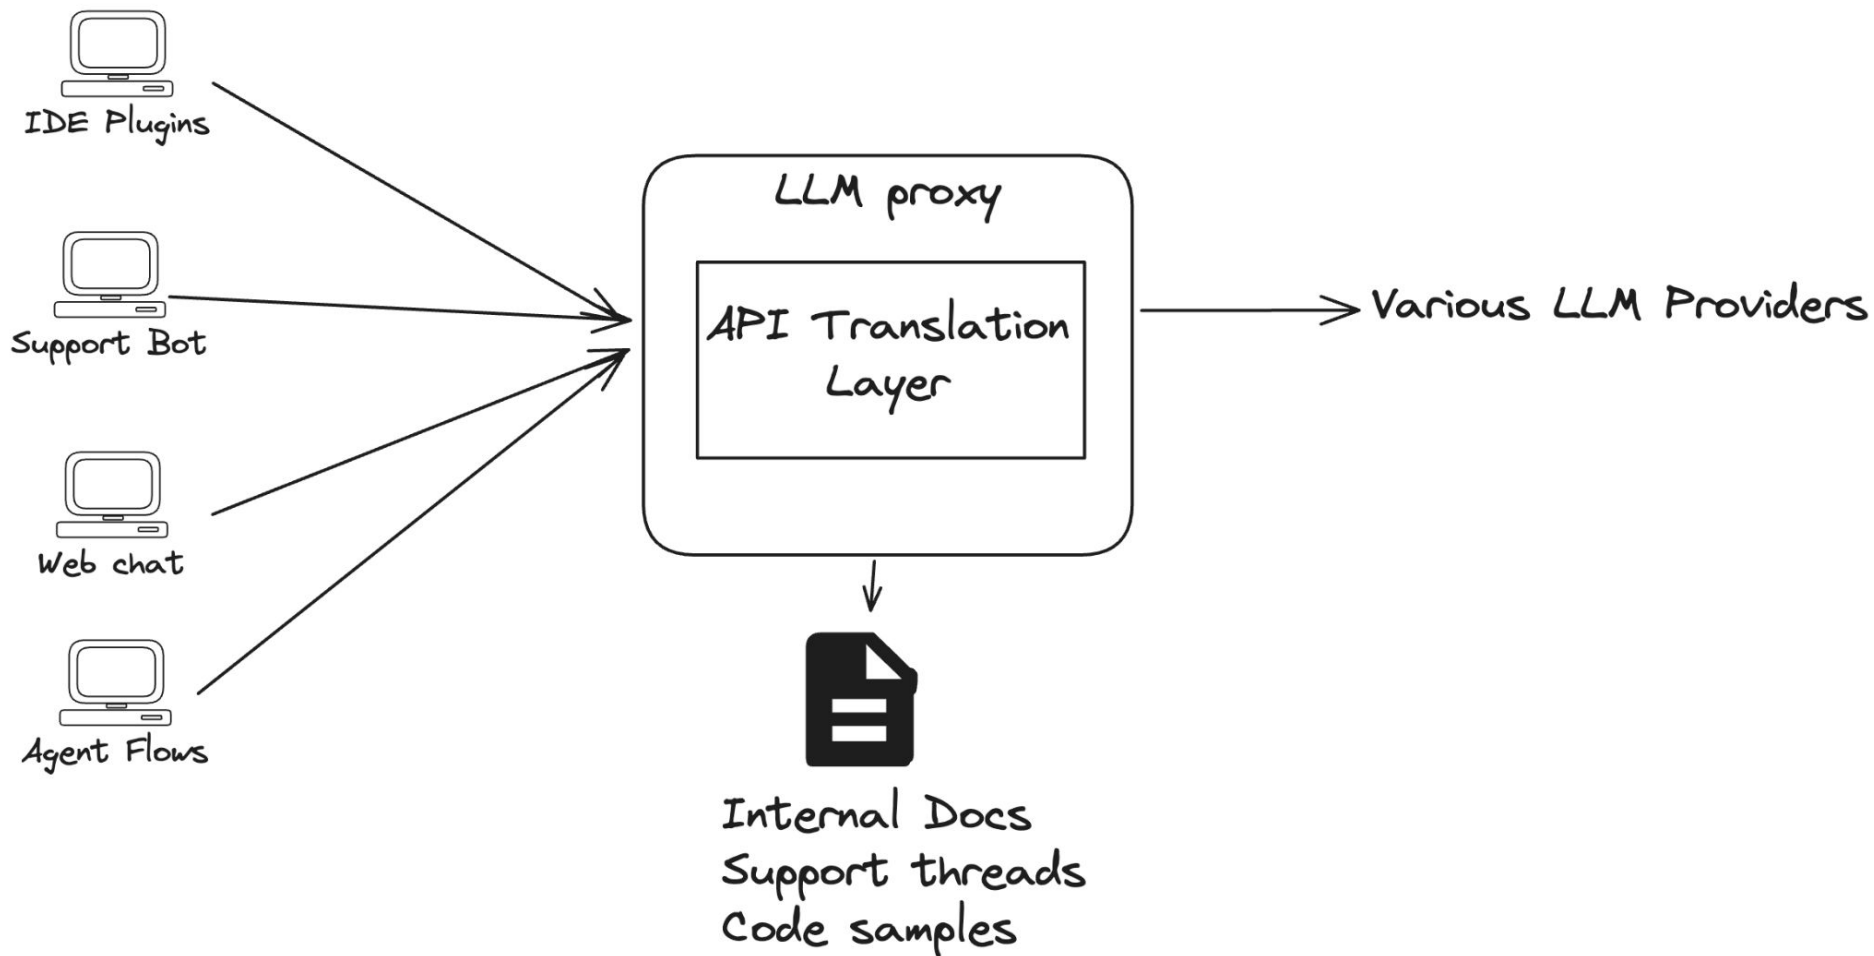
Platform for Rapid Prototyping with LLMs

LLM proxies

SDKs in various languages

“Sandcastle”: Service prototyping framework

Evaluation Framework



Review GPT

What if AI reviewed code?

Tag AI reviewer



AI review



Conversation 2 ← Commits 1 Checks 2 Files changed 1

commented 3 hours ago · edited

Change summary:

Prewarm simulator to avoid xcodebuild timeouts when trying to launch the simulator.

How was it tested?

Please review:

@svc-review-gpt

Prewarm simulator 733e8eb

deployboard commented 3 hours ago

🌟 **ReviewGPT AI-Generated Code Review** 🌟

1. While this PR seems to make sense conceptually, it's also worth considering that Apple's Simulator app may not be running by the opening of the Simulator before the tests since killing the Simulator processes is part of the test workflow. The opening of the simulator might be ignored, or the workspace may have some unexpected state when the tests run. Moreover, the `open -b com.apple.iphonesimulator || true` command runs asynchronously, and there is no certainty the simulator app will be ready when the tests start running.

I propose a more robust way handling simulator prewarming by utilizing `simctl` command which provides direct capabilities for interacting with iOS simulators. For example, you can directly boot a simulator device using `xcrun simctl boot <device>`, where `<device>` is the id of the simulator, and you can ensure this operation completes before running the tests.

Here is my recommendation:

```
system_and_log 'pgrep Simulator | xargs kill || true'
```

```
# Prewarm simulator to avoid xcodebuild timeouts when trying to launch the simulator.  
- system_and_log 'open -b com.apple.iphonesimulator || true'  
+ system_and_log 'xcrun simctl list devices | grep "Shut Down" | cut -d "(" -f2 | cut -d ")" -f1 | xargs -I {}
```

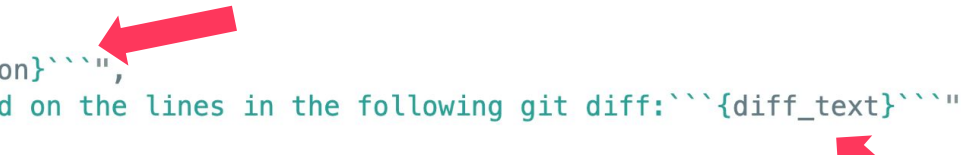
This command lists all devices, filters the ones that are shut down, extracts the id, and boots them.

Note: This command attempts to boot all shut down simulators. If you want to boot a specific simulator, you would need to identify it by its identifier instead of the `grep "Shut Down"` command.

Review GPT

User prompt

```
# It's important that any dynamic input sources are truncated here (such as git diff output)
# otherwise, the API response will fail if the number of tokens exceeds the allowed threshold.
prompt_text = "\n".join([
    f"PR title:``{pr_title}``",
    f"PR description:``{pr_description}``",
    f"Provide useful suggestions based on the lines in the following git diff:``{diff_text}``",
])
```

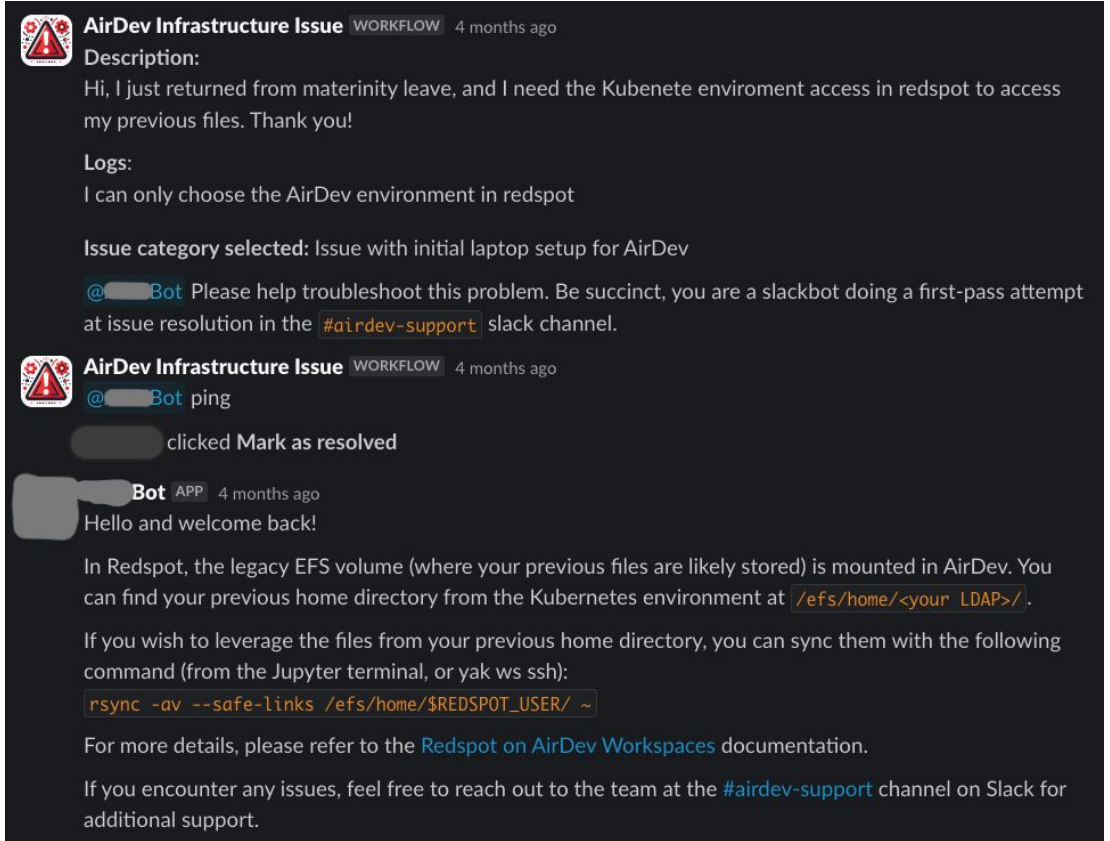


System prompt

```
def system_prompt():
    return "\n".join([
        f"Assistant is {REVIEW_GPT_GITHUB_USER} aka \"ReviewGPT\", an AI that does code reviews of Airbnb-internal code in Github.",
        "Assistant will review the git diff, where lines starting with '-' are being removed, lines starting with '+' are being added, and other lines are surrounding context.",
        "Assistant does not introduce itself. Assistant's responses ONLY include the actual code review without any other context.",
        "Assistant ONLY uses backticks `` to surround code snippets and never uses backticks to format non-code.",
        "Assistant does not criticize for minor style issues and does not criticize for things that are not in the diff, like missing PR description items.",
        "Assistant offers concrete, actionable, and numbered suggestions based on best practices and observed bugs, and does not suggest tangential, vague, nitpicky, or out-of-scope changes.",
        "If Assistant doesn't have any constructive review comments to add, it will simply respond 'LGTM'.",
        "100% of the response is always in github markdown (without any other text or headings, and without escaping the markdown), suitable for copying into a github comment directly.",
    ])
```

AI Support Bot

Can AI respond to engineering support requests?



AirDev Infrastructure Issue WORKFLOW 4 months ago

Description:
Hi, I just returned from maternity leave, and I need the Kubenete enviroment access in redspot to access my previous files. Thank you!

Logs:
I can only choose the AirDev environment in redspot

Issue category selected: Issue with initial laptop setup for AirDev

@**Bot** Please help troubleshoot this problem. Be succinct, you are a slackbot doing a first-pass attempt at issue resolution in the [#airdev-support](#) slack channel.

AirDev Infrastructure Issue WORKFLOW 4 months ago

@**Bot** ping

Bot clicked Mark as resolved

Bot APP 4 months ago
Hello and welcome back!

In Redspot, the legacy EFS volume (where your previous files are likely stored) is mounted in AirDev. You can find your previous home directory from the Kubernetes environment at `/efs/home/<your LDAP>/`.

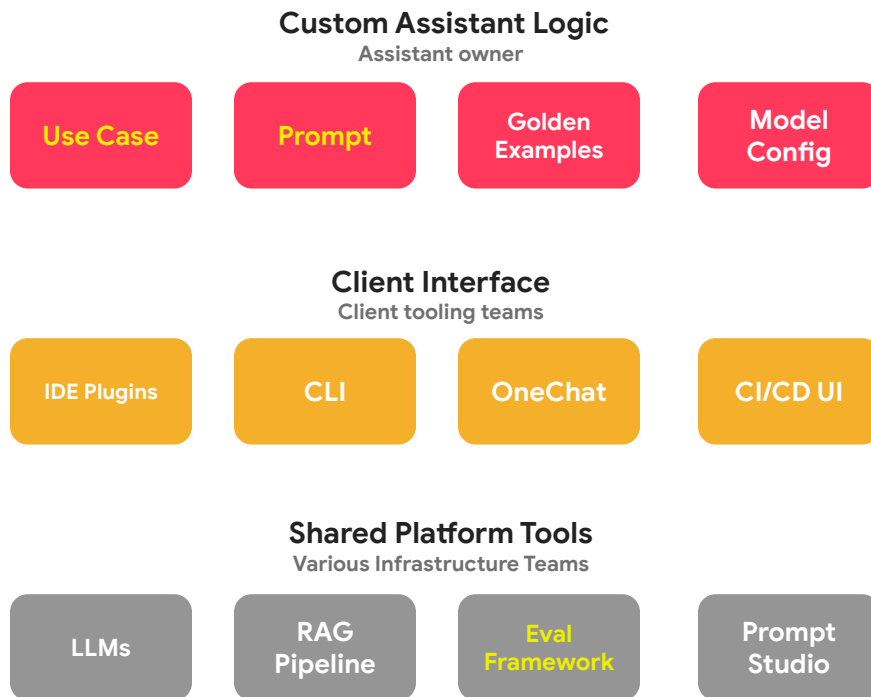
If you wish to leverage the files from your previous home directory, you can sync them with the following command (from the Jupyter terminal, or yak ws ssh):

```
rsync -av --safe-links /efs/home/$REDSPOT_USER/ ~
```

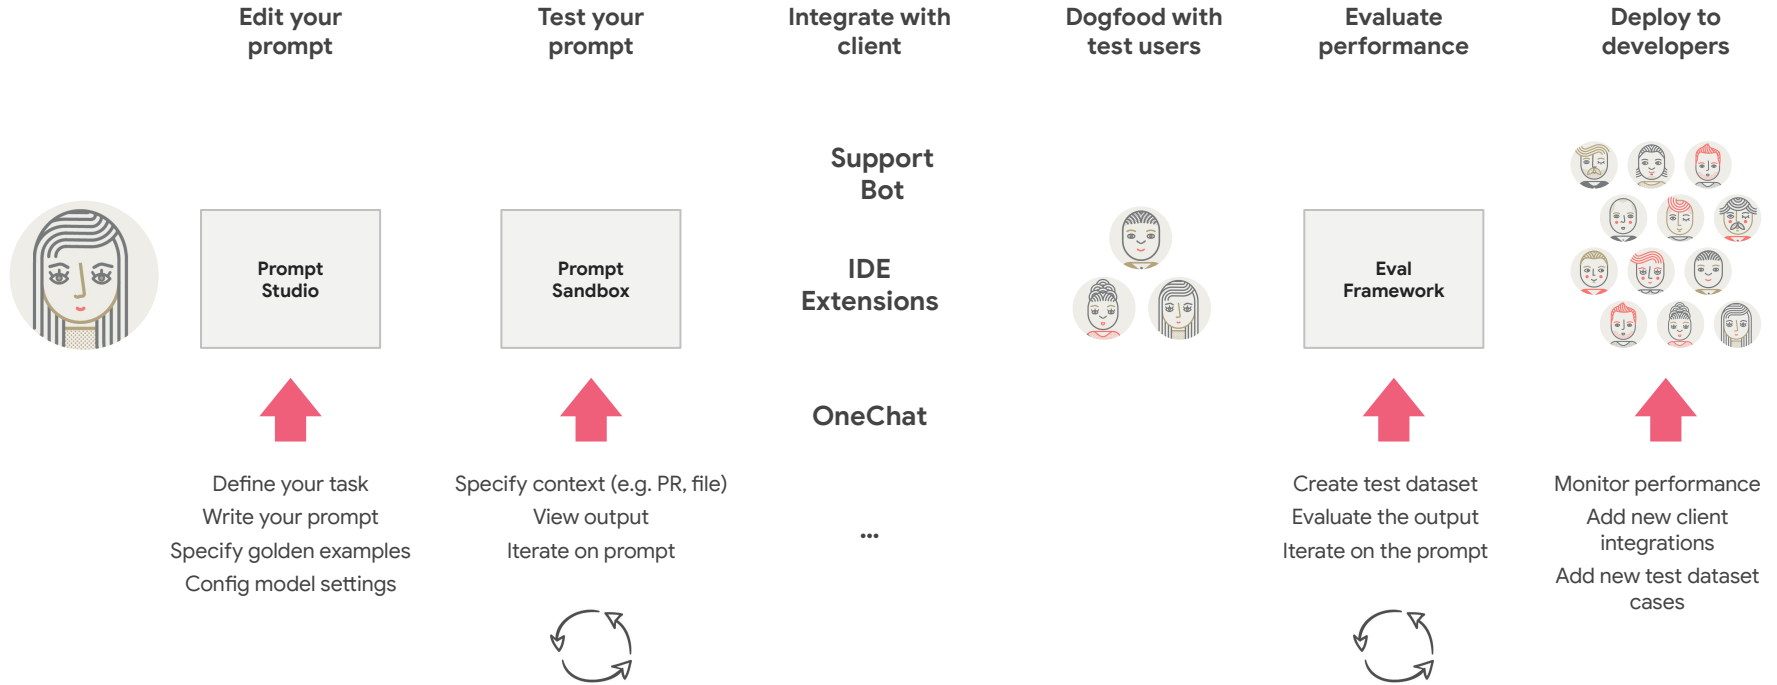
For more details, please refer to the [Redspot on AirDev Workspaces](#) documentation.

If you encounter any issues, feel free to reach out to the team at the [#airdev-support](#) channel on Slack for additional support.

Tightly-scoped Assistant Architecture



Create Your Own Assistant



Thank you!

Questions?

