



Your LLMOps journey with CircleCI

Hello there,

We're excited to share the new tools and features to support the builders of AI-enabled applications. Here are the two updates we've rolled out this week:

1 - [CircleCI Evals Orb](#)

This new orb makes it easy to integrate LLM evaluations into a CI pipeline, and to access and review evaluation results without context switching. It provides commands to store the evaluation outputs as artifacts and also to add them as comments to PRs.

Currently, the [Evals Orb](#) exposes commands to run evaluations through [LangSmith](#) and [Braintrust](#). If your evaluations leverage a different tool, [let us know!](#) You can test out this orb by forking the [llm-eval-examples](#) sample repo.

2 - LLMOps credential management

We made it easier to store OpenAI credentials as contexts in CircleCI. The new LLMOps integration page lets you enter, verify and access your OpenAI credentials, as well as providing a starting template for your config.yml file.

[Navigate to Project Settings > LLMOps to get started](#)

Once this is done, you can also save the credentials to your preferred evaluation platform in the Advanced section of this same LLMOps Settings window.

For now, [Braintrust](#) and [LangSmith](#) are available, but we will soon provide more flexibility.

If you are interested in running evaluations on CircleCI but have questions or concerns, simply answer this email (ai-feedback@circleci.com) or [book a chat with our team!](#)

To better building,
Alex and the CircleCI for AI team

Useful resources

- [Docs - Evaluating LLM-enabled applications with CircleCI](#)
- [Course - Automated Testing for LLMOps](#)

