



Select Committee on Adopting Artificial Intelligence (AI)
Parliament House
Canberra ACT 2600
aicommittee.sen@aph.gov.au

23 August 2024

Atlassian is grateful to the Committee for the opportunity to appear before the Senate Select Committee on adopting Artificial Intelligence (AI) on 17 July 2024.

On 9 August 2024, Atlassian received the below questions on notice from the Chair of the Committee. Noting that Atlassian did not commit to take these questions on notice during its hearing, we remain supportive of the Committee's inquiry. As such, we welcome the Committee's ongoing interest in these issues and have provided answers to the Chair's questions below on a voluntary basis and subject to Atlassian's confidentiality commitments (including with its third party partners).

Question 1

Stanford University's Centre for Research on Foundation Models produces a regular [report](#) issuing transparency scorecards to the developers of Large Language Models, including OpenAI. It assesses their performance against 100 indicators of transparency. The most recent analysis, released in May of this year, awarded OpenAI a transparency score of just 49%. It said their data sources are not disclosed, neither are the creators or curators of the data, nor what the selection protocols for the data are.

- a. Was Atlassian given any of that information as part of its agreement with OpenAI?*
- b. And if not, why does Atlassian feel comfortable integrating OpenAI models in its products if it doesn't know the most basic information about how it works?*

Atlassian and OpenAI first entered into, and [announced](#), their partnership in early 2023.¹

Since that time, the market for large language models (**LLMs**) and other foundation models has made incredible strides in terms of quality, performance and transparency. The Stanford Center for Research on Foundation Models' [Foundation Model Transparency Index](#) referred to in this Question also makes this clear, emphasising many of these changes and improvements from the first Index in October 2023 to the revised Index in May 2024.² We expect these increases to continue.

While the details of Atlassian's contractual agreements with OpenAI are commercial in confidence, we confirm that we are always looking to the developers of the LLMs that we use to provide the necessary information and transparency about their models, and where and how they perform best. This includes several of the information sources noted in the CFRM Index, such as detailed model and system cards.³

We also continually assess, test and evaluate LLMs ourselves across a range of factors — combining public benchmarks, our own quality evaluations and other factors — to confirm their suitability to the use cases of our products, and do not solely rely on the information provided by developers.

¹ See, e.g., <https://www.cnbc.com/2023/04/19/atlassian-taps-openai-for-atlassian-intelligence-generative-ai-launch.html>.

² See <https://crfm.stanford.edu/fmti/May-2024/index.html>.

³ See, e.g., the GPT-4o system card released by OpenAI: <https://openai.com/index/gpt-4o-system-card/>.

Further, this assessment is not limited to OpenAI's models. To provide the best offering to our customers, we want the best and most appropriate models for our specific use cases and customer needs, and our teams are continually exploring and assessing a range of LLMs accordingly.

Question 2

The Committee has heard from a number of organisations who have expressed concerns about the legitimacy and accuracy of the information generated by LLMs, especially when they experience 'hallucinations.'

a. Given Atlassian's partnership with OpenAI, what oversight do you have over the accuracy of OpenAI's outputs in order for you to integrate it in the products and services you provide to your customers?

b. When there are inaccuracies in the information provided by OpenAI and then onto your customers, what is the process for you to determine if the liability sits with Atlassian or OpenAI?

Please refer to our response to Question 1 above with respect to our partnership with OpenAI.

When it comes to output quality and accuracy, we are always working to improve the quality of the responses that our customers and users receive when they use our AI-powered products and features.

We are also clear with our customers and users, including in our customer-facing documentation, [transparency materials](#) and in-customer experiences,⁴ that many of the generative AI models used as part of those features provide responses based on users' inputs and are probabilistic in nature. As a result, these models can sometimes behave in ways that are inaccurate, incomplete, or unreliable.

We encourage, and prompt, users to think about the situations when they use Atlassian Intelligence and review the quality of the responses they receive before sharing them with others.

⁴ See, e.g., <https://www.atlassian.com/trust/atlassian-intelligence> and <https://www.atlassian.com/trust/atlassian-intelligence/transparency#introduction>. For more on Atlassian Intelligence, see <https://www.atlassian.com/platform/artificial-intelligence>.

