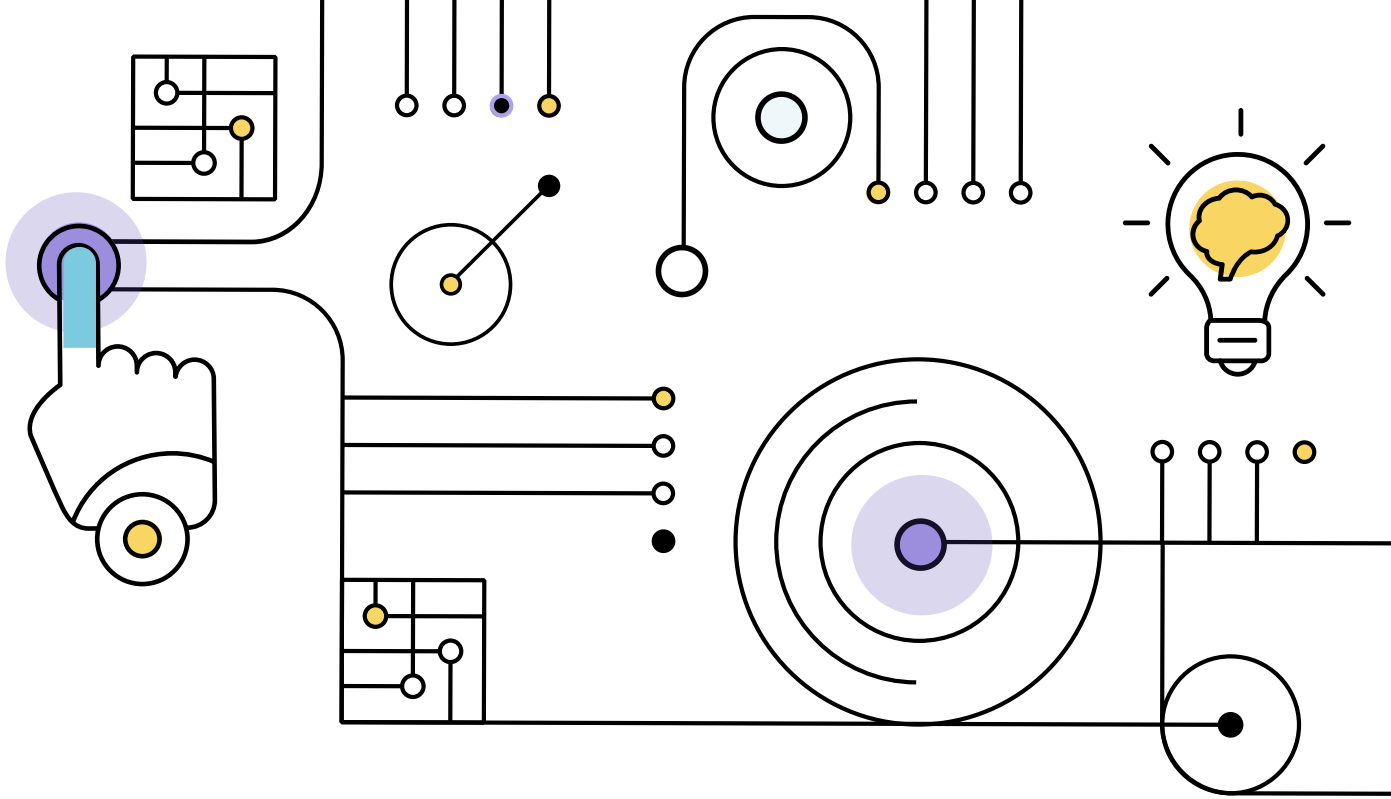

WorkBoard Gen AI Data Security Questions

March 2024

WorkBoard

FAQ

How does the OKR functionality in WorkBoard utilize generative AI?	WorkBoard leverages Azure Managed OpenAI and Anthropic Claude Services to provide intelligent suggestions and prompts for generating OKRs based on the data available in the platform.
What steps are taken to ensure the security and privacy of our OKR data when using this functionality?	At WorkBoard, we take the security and privacy of your data seriously. We follow industry best practices and employ robust security measures to protect your OKR data. All communication between WorkBoard, Azure OpenAI and Anthropic Claude is encrypted, and access to your data is strictly controlled and limited to authorized personnel.
How is the OKR data transmitted to Azure OpenAI and Anthropic Claude, and what measures are in place to secure this data during transmission?	The OKR data is securely transmitted to Azure OpenAI and Anthropic Claude using encrypted channels and industry-standard security protocols. We adhere to secure data transmission practices to ensure the confidentiality and integrity of your OKR data.
Are there any measures in place to prevent unauthorized access or misuse of our OKR data within Azure OpenAI and Anthropic Claude?	Yes, Azure OpenAI and Anthropic Claude maintains a robust security infrastructure to prevent unauthorized access or misuse of customer data. They have implemented strict access controls, monitoring systems, and auditing mechanisms to ensure the protection of your OKR data.
What level of control do we have over our OKR data?	You retain full ownership and control over your OKR data. Azure OpenAI and Anthropic Claude does not use or retain your data beyond the scope of the OKR generation process. WorkBoard does not share your data with any third parties without your explicit consent.
Are there any data anonymization techniques used when transmitting our OKR data to Azure OpenAI and Anthropic Claude?	Your data is encrypted in transit. Personally identifiable information (PII) or any sensitive information is anonymized or tokenized before being sent to Azure OpenAI and Anthropic Claude, ensuring the privacy and confidentiality of your data. See architecture diagram on following slide.
Can we trust the suggestions and prompts generated to maintain the confidentiality of our OKR data?	Yes, the WorkBoard Co-Author is designed to provide high-quality suggestions while maintaining the confidentiality of your OKR data. The model focuses on generating relevant and useful prompts without retaining or disclosing any sensitive information.
Does Azure OpenAI and Anthropic Claude have any data retention policies regarding our OKR data?	Azure OpenAI does not retain your OKR data beyond the scope of the OKR generation process. Once the suggestions are generated, the data is discarded, ensuring that your data remains secure and private.
Are there any additional security certifications or compliance standards that Azure OpenAI and Anthropic Claude adheres to?	Azure OpenAI maintains various security certifications and compliance standards, including ISO 27001, SOC 2 Type II, TISAX and GDPR. These certifications ensure that their infrastructure and practices meet the highest security and privacy standards.

Architecture diagram

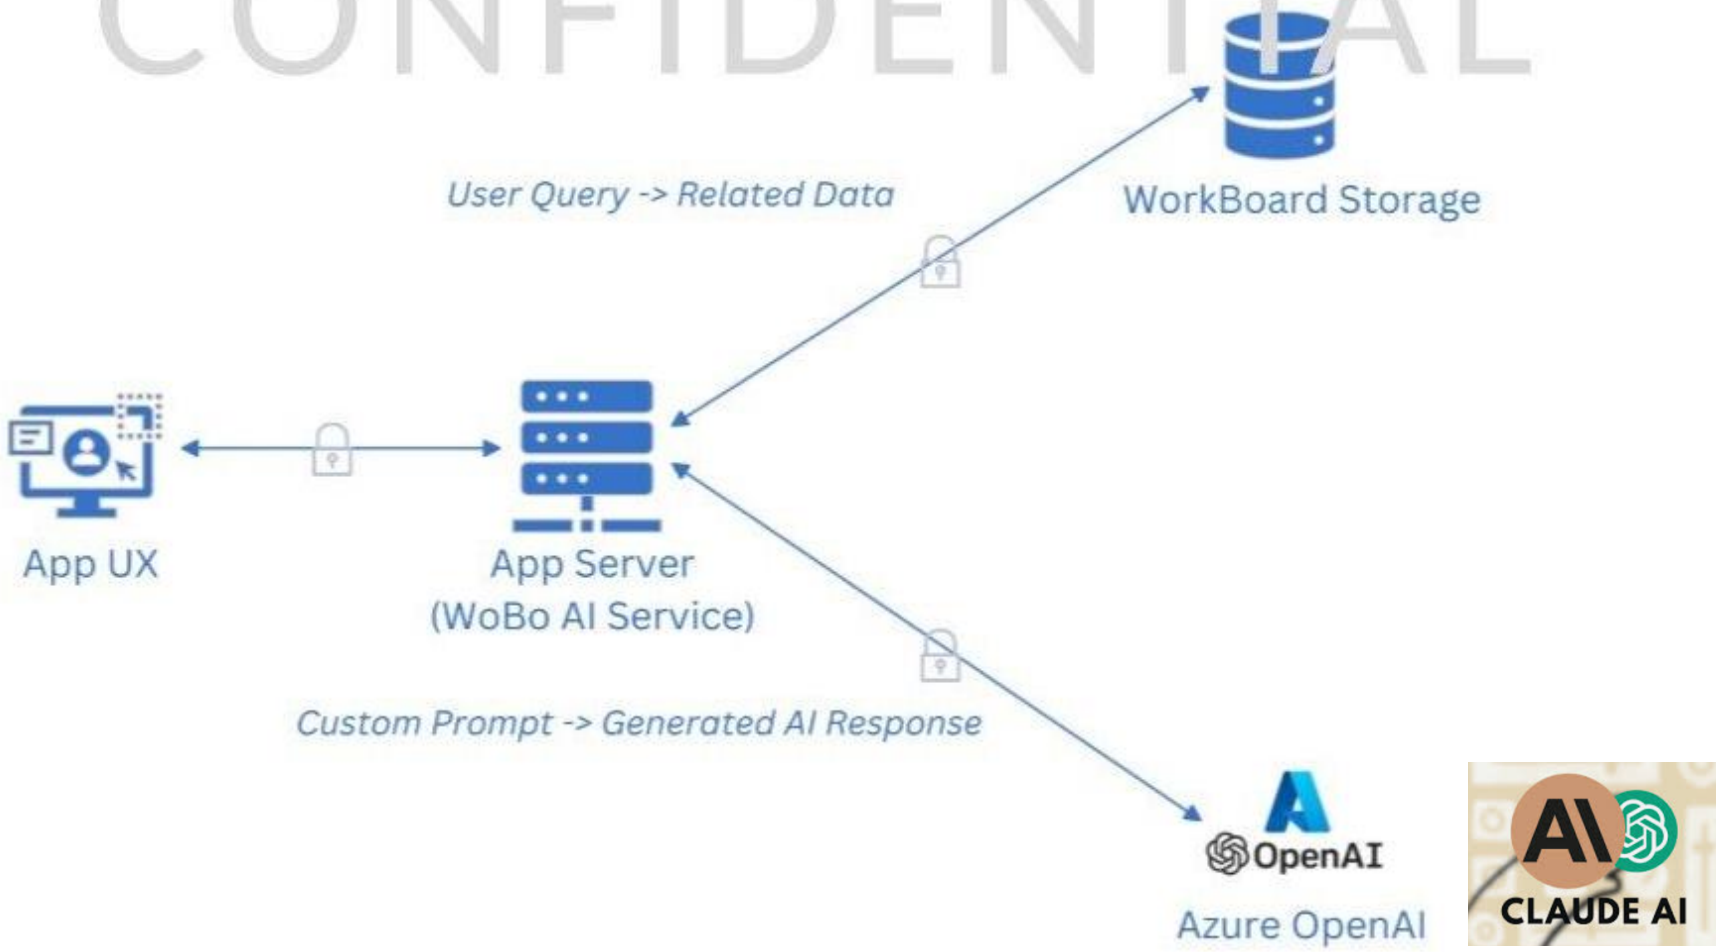

WorkBoard's Generative AI

General questions

Is access to OpenAI via their public API?	No. We use Azure OpenAI. Microsoft hosts these models, and the model we use is exclusively used for WorkBoard. More details in the blue highlighted section here: https://learn.microsoft.com/en-us/legal/cognitive-services/openai/data-privacy?context=%2Fazure%2Fai-services%2Fopenai%2Fcontext%2Fcontext
Would customer data/content flow to OpenAI in a way that would give them a right to use customer data/content?	No
Has WorkBoard established a private usage service which they own and control using, e.g. a private Azure Enterprise service to which they subscribe?	Yes
If a private usage service has been established, what claim(s) would WorkBoard expect to have over their right to use customer data / content?	It's based on our contract with Azure and the strict policy established by Azure. See FAQ.
Can you provide an overview of the specific AI technologies which are used by WorkBoard Co-Author, including its processes, functions, user interfaces and outputs?	We use input prompts, which are built with WorkBoard data, accessible by the authorized user to generate customized responses. We use Generative AI models hosted in Azure OpenAI exclusively deployed for WorkBoard.

Training data and data residency

Can you describe the data sources used to train the AI solution, and how the solution is trained on these data sources?	We do not use data for training any model.
What input data is required from the customer/user?	Team names, previous OKRs and upline manager's OKRs are automatically fetched from WorkBoard sources. In some cases, users have the option to add their team's priorities in the form of text input.
How can you demonstrate that the data used to train the AI is diverse, free of bias and representative?	To ensure the diversity and representativeness of the data used to train our AI, we adhere to stringent guidelines and methodologies. We leverage Azure's and Antrhopic's recommended practices to curate and process our datasets, striving to make them as unbiased and inclusive as possible. It's important to note that while we've taken measures to reduce bias in the AI, its responses are also shaped by user input. For a comprehensive overview of our data transparency, including how we handle and mitigate potential biases, please refer to the detailed note provided by Microsoft: https://learn.microsoft.com/en-us/legal/cognitive-services/openai/transparency-note?tabs=text
Where will data shared by customer/users with the AI solution reside? Will data shared by customers/users be transferred outside of the United States, United Kingdom or European Economic Area?	The data won't be shared or used outside of WorkBoard ecosystem. See FAQ.
Will customer's/user's use of the AI solution involve the collection and processing of personal data by the AI solution?	We do not use any personal data for generating OKRs.
What retention periods are in place for data shared with the AI solution, and what is the process for how that data will be securely deleted once no longer required?	Data won't be stored outside of the Workboard ecosystem. In addition to the existing policies, we opted out of optional abuse monitoring. More details here: https://learn.microsoft.com/en-us/legal/cognitive-services/openai/data-privacy?context=%2Fazure%2Fai-services%2Fopenai%2Fcontext%2Fcontext#how[...]iew

Audit and compliance

Can you describe your internal audit process which provides assurance that the AI solution complies will all relevant laws and regulations – particularly should these change over time?	We maintain a dynamic internal audit process to ensure our AI solutions comply with all relevant laws and regulations. Our compliance team continuously monitors regulatory updates, while our AI systems are reviewed to detect deviations. Periodic internal reviews, coupled with annual external audits, assess our adherence for e.g. PEN tests. Each development phase is thoroughly documented, and findings from audits are rapidly integrated back into our systems and processes. Regular training ensures our team stays updated, and our process's iterative nature ensures swift adaptation to regulatory changes.
How would you be able to give customers oversight on the AI's data engineering, data science and MLOps should we require this?	We have a dedicated Data Science and AI/ML team, who can work with customers to provide oversight needed to ensure they have a comprehensive view of our AI development processes.

Accuracy

How are the solution's outputs and conclusions are validated?	The use of AI in OKR Co-Author generates suggestions for Objectives and Key Results that users can accept, further modify and use AI to refine. The implementation is designed to inspire options and help teams as they brainstorm OKRs, but users need to make the final decision of what they select and capture as their OKRs. With this in mind, our solution's OKR outputs undergo a multi-layered validation process during the design process of the solution, post deployment and every time users engage with the solution. During the design of the solution, we employ cross-validation techniques to compare model outputs with WorkBoard best-in-class recommendations of Objectives and Key Results during the model selection and prompt tuning to ensure validity and usefulness of suggestions. Post-deployment, periodically, our team conducts manual audits and reviews of the outputs. Furthermore, the design of Co-Author to generate suggestions that users can leverage or dismiss is a continuous feedback loop. Quantitative data from acceptance rate of suggestions, in addition to qualitative data from beta and early release cycles will also trigger reviews and refinement of the underlying algorithms to maintain the highest standards of relevance for suggestions
What testing has taken place to remove or mitigate against the risk of "hallucinations" or other inaccuracies in the outputs created by the AI solution?	WorkBoard integrates context into the prompts to combat hallucinations. This context comprises the team's previous OKRs, upline OKRs, and other team-level data stored in WorkBoard. Furthermore, users typically provide additional context, which aids in generating more accurate responses. Our AI/ML engineers at WorkBoard continuously monitor the system and adjust the AI. We do not send any data to OpenAI temperature to further reduce hallucination.

Intellectual property and third-party licenses

Who is the owner of the intellectual property in the AI solution?	WorkBoard Inc.
Has the AI been trained on data which is owned or produced by 3rd parties? If it has, have those 3rd parties given consent to their data being used in this way?	WorkBoard uses the GPT3.5 as the foundational LLM that has been trained from public data by OpenAI and hosted in a private Azure instance by Microsoft
Is there any ongoing litigation (in any jurisdiction in the world) involving the lawful use of the AI solution, or any litigation reasonably in prospect?	We have no pending or anticipated litigation on any matter related to our AI solution.

Deployment

Can you describe the deployment process? What deployment methods do you support: batch, API, on the edge?	Our primary deployment method revolves around APIs, enabling real-time interactions and seamless integration with existing systems. While we currently don't employ batch or edge deployments, we enhance our API deployments with feature flagging. This allows us to manage new feature rollouts incrementally, ensuring stability, and facilitating immediate feedback, optimizing the user experience and system reliability.
Do you use MLOps methodology? Please describe how this continually improves the process from data sourcing and labelling through to deployment and monitoring?	WorkBoard uses Phase 1 and 2 of MLOps methodology and a subset of Phase 3, specifically - testing, manual versioning and monitoring. We are always evolving our processes, and automating steps, with the goal of maturing our MLOps.
What user support will be required/provided? How will issues be reported, investigated and resolved?	To support our AI/ML users and simplify interactions with our tools and team, we utilize the same support channel and funnel that we currently employ for other product areas and issue tracking.
Does the solution depend on integration with any other systems – if so, what are these and how are they integrated?	WorkBoard's AI/ML solution doesn't depend on the integration of the other system.

Gen AI in development

What are the guardrails for WorkBoard's use of GitHub Copilot and other code-complete tools?	We recognize the importance of maintaining a robust and secure development environment. Where Gen-AI tools are utilized in our development practices, they have undergone a thorough security review with our security team. We enforce stringent guardrails around AI-generated solutions, including regular code and pull request reviews, adherence to licensing constraints, and rigorous security assessments, ensuring our codebase remains secure and compliant with our standards.