betterworks

Questions & Answers for HR AI Software Vendors

If you are considering adopting a generative AI solution as part of your HR tech stack, these questions will help you and your IT and information security teams determine if the solution is secure, compliant, and ethically sound. Understanding data privacy, model training, and vendor practices helps mitigate risks associated with sensitive employee information and prevents potential legal issues. Asking vendors to respond to these questions will help you make an informed decision that aligns with the organization's values and regulatory requirements.

Data Privacy and Security

- Where is our employees' data stored? Is it encrypted both in transit and at rest?
  - **Answer:** All employees data is stored securely within the our cloud solution deployed on Amazon Web Services (AWS) cloud infrastructure. AWS offers a robust suite of security features and adheres to strict compliance standards. All employee data is encrypted at rest using the industry-standard AES-256 encryption algorithm. Encryption keys are managed using AWS KMS. This significantly reduces the risk of unauthorized access to the data, even in the unlikely event of a breach. Data is encrypted in transit using TLS 1.2+ with security certificates leveraging asymmetric encryption & digital signatures (Elliptical Curve & RSA). SHA2 is used for hashing. Betterworks also provides customers with an SFTP for uploading data. This leverages SSH.
- What security measures are in place to protect employee data from unauthorized access?
  - **Answer:** We employ a multi-layered approach to security, including: strict access controls within our AWS environment based on the principle of least privilege. We also have firewalls and intrusion detection/prevention systems (IDS/IPS) to monitor network traffic and identify and block any suspicious activity that might indicate unauthorized access attempts. We conduct regular penetration testing and vulnerability assessments of our AWS environment to proactively identify and address any potential security weaknesses. These measures are audited and verified by external third-parties under our SOC-2 Compliance and ISO 27001 certification. More can be found here: **https://www.betterworks.com/privacy-notice/**

- Does your solution comply with relevant data privacy regulations (e.g., GDPR, CCPA)?
  - **Answer:** Yes, Betterworks is in compliance with the General Data Protection Regulation (GDPR) in the European Union, the California Consumer Privacy Act (CCPA) in the United States, the EU-U.S. Data Privacy Framework (EU-U.S. DPF) program, the UK Extension to the EU-US DPF, and the Swiss-U.S. Data Privacy Framework (Swiss-U.S. DPF) program as set forth by the US Department of Commerce (DOC).

- How do you handle data retention and deletion policies, especially concerning generated content?
  - **Answer:** We retain employee data for as long as the employee is active with your company and our solution is in use. This allows us to provide the best possible service. The customer has the right to request deletion of employee data at any time. Once you submit a deletion request, we will remove the data from our systems within (30) days. Generated content is treated the same way as all other user data. This means it follows the same retention and deletion policies outlined above.

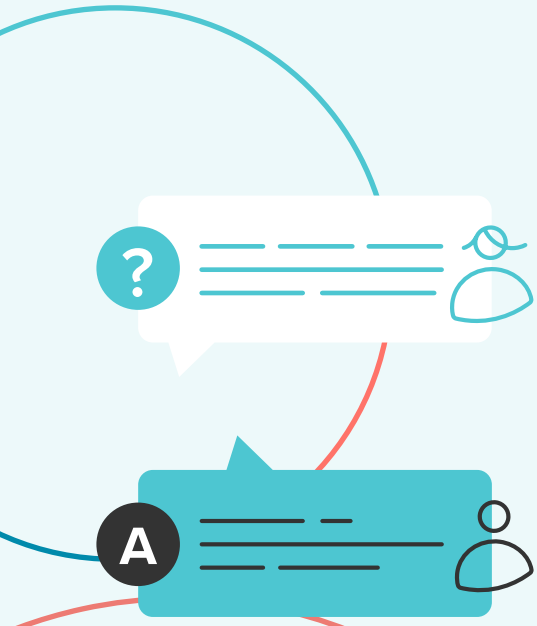

Model Training and Data Sources:

- Is our employees' data used to train AI models?
  - **Answer:** No, your employees' data is never used to train or fine-tune our generative AI model. We understand the importance of data privacy and take steps to ensure your employees' data is used solely for the purposes of your specific use case with our solution.

- Are you fine-tuning your generative AI models? If yes, what data sources are used to achieve that?
  - **Answer:** Currently, we don't perform fine-tuning on your employee data. However, we are constantly exploring ways to improve the quality and accuracy of our generative AI models. In the future, we plan to leverage synthetic data for fine-tuning.

- Are there any steps taken to mitigate bias in the training data and model outputs?
  - **Answer:** We leverage prompt engineering to guide our AI models towards generating unbiased outputs. By carefully crafting prompts, we can steer the model away from generating outputs that reflect biases. We also continuously monitor our model's performance to identify potential bias in its outputs. We are introducing more safeguards and guardrails over time to ensure we capture any biases in the inputs or outputs.

- How do you ensure the model is not generating outputs that could be discriminatory or offensive?
  - **Answer:** Similar to bias prevention we use identical prompt engineering techniques to control inputs and outputs for any kind of misuse.

- How frequently are the models updated and retrained?
  - **Answer:** We understand the importance of keeping our models up-to-date to ensure optimal performance and address potential issues. While we don't have a pre-defined schedule for updates, we are committed to forward compatibility. Our system is designed to integrate future model updates seamlessly, minimizing disruption to the user experience.

Algorithm Transparency and Explainability:

- Can you explain how your generative AI models make decisions or create content?
  - **Answer:** At Betterworks we are committed to transparency and explainability. Every use case of generative AI has a section to explain what happened and why. Our aim is to empower our users to not only leverage the power of generative AI but also understand the reasoning behind its outputs.

- How do you handle false positives/negatives or errors in AI-generated content?
  - **Answer:** Unlike chatbots that answer open ended questions and can potentially provide inaccurate or misleading information, our AI is focused on specific tasks. We tailor it for tasks like summarizing large amounts of data, finding relevant information within a dataset, and generating content based on existing information you provide. This targeted approach helps minimize errors by ensuring the AI is working with well-defined parameters. We also have quality control measures in place to identify and address any potential inaccuracies or misleading outputs before they are shown to the user.

Vendor Security Practices & Legal Considerations:

- What are your policies regarding third-party access and subcontractors related to generative AI?
  - **Answer:** We restrict access to your employee data to authorized personnel within our company and any subcontractors directly involved in operating the generative AI solution. In the use case of generative AI in particular, we use OpenAI as a sub-processor and they are contractually obligated to uphold data security standards that meet or exceed our own. See more here: **https://www.betterworks.com/ai-terms-and-conditions/**

- Is your generative AI model based on proprietary technology or a third-party vendor (like OpenAI)?
  - **Answer:** We initially leveraged OpenAI's GPT models through their Enterprise API for a proven solution. However, we're also deploying mature enterprise-grade open-source LLM models into our secure AWS environment. This gives our customers the option to choose between a well-established solution and a self-deployed model that can leverage the scalability and security of our AWS cloud environment.

- How do you ensure compliance with industry-specific regulations concerning generative AI?
  - **Answer:** While specific regulations for generative AI are still emerging, we prioritize responsible AI use. We work closely with each customer to understand their industry's potential regulations and adapt our practices to align with the latest developments. We focus on transparency, conduct risk assessments, and actively monitor the evolving landscape to ensure responsible AI development alongside our customers.

- What legal protections and liabilities are in place if the AI system generates harmful or inappropriate content or if breaches occur?
  - **Answer:** While our focus is on empowering you with control over the AI's outputs, we understand potential concerns. Disclaimers in our contract limit our liability for content accuracy, and robust security measures minimize data breach risks. However, you have ultimate control over generated content and we recommend consulting your legal team to understand specific liabilities in your situation. Read more here: **https://www.betterworks.com/ai-terms-and-conditions/**