ACLU– Generative AI Vendor Questionnaire

Last updated: October 2025

This questionnaire seeks to gain some insight into your company's policies and practices around generative AI systems and related security, privacy, copyright, bias, fairness, transparency, and environmental impact issues. **This questionnaire should be completed only if 1) the service(s) or solution in question has generative AI capabilities, and 2) those capabilities cannot be "turned off" or disabled at the root level.**

This questionnaire is intended to supplement the ACLU's security and privacy reviews, which generally apply both to systems that use generative AI or that do not use generative AI or whose generative AI capabilities can be "turned off." **This is not a comprehensive list of all security and privacy questions the ACLU asks during procurement – only those that relate specifically to Generative AI.** The questions included reflect the American Civil Liberties Union and American Civil Liberties Union Foundation's ethical AI principles and practices.

Company Name:		
Person Completing Questionnaire	Name:	
	Title or Role:	
Description of the services or solution the company is providing to ACLU:		
2FA integrated app?		
Date:		

[Prerequisite to be completed by ACLU staff member]

Question	Yes	No	Additional Comments or Clarifications	For ACLU Use Only
Are you interested in procuring this tool in part for its generative AI capabilities? If so, use the additional comments section to describe ACLU intended or potential uses of the system's generative AI capabilities, including both intended immediate/short-term uses and possible longer-term uses.				

Instructions

1. Provide a response for each item on the questionnaire after assessing the adequacy and effectiveness of the documentation, testing, or control procedures within your area.

2. For each question; please indicate an "x" under the columns for Yes, No, or N/A.

    2.1. In the "Comment" section enter or attach a description of the control process(es) or any information that supports or clarifies your responses.

ACLU– Generative AI Vendor Questionnaire

ACLU

2.2. Indicate what evidence exists to support your responses or cross-reference to the supporting documentation.

2.3. Indicate 'Not Applicable' (N/A) for questions that do not apply to your systems and provide explanation in the "Comments" section.

3. For narrative answers, please provide explanations in the "Comments" section.

4. Return the completed questionnaire to the ACLU Generative AI Working Group.

Question	Yes	No	N/A	Additional Comments or Clarifications	For ACLU Use Only
Part 1: General Transparency					
1.1. Does the product incorporate AI that is not generative AI? If yes, please briefly summarize.					
1.2. If you answered yes to 1.1, explain how your company distinguishes between non-generative AI and generative AI in the context of your product.					
1.3. Is there a model card, datasheet, data statement, system card, or other data and model documentation available for the system(s) in question? If so, please provide copies. If not, please explain why not.					
1.4. If the documentation described in Question 1.3. is available, has it been made available to the public? If not, why not?					
1.5. If documentation described in Question 1.3 is not available, please provide a description of the system(s) in question, including information about how the system(s) incorporates or relies on generative AI systems, who developed the system(s) and its components, how the system is intended to be used and how it should not be used, and related documentation.					
1.6. Please summarize and provide copies of all available documentation related to the system(s)'s accuracy and other aspects of model performance, including a description of what metric(s) were chosen for measuring accuracy.					

Question	Yes	No	N/A	Additional Comments or Clarifications	For ACLU Use Only
1.7. Please summarize and provide copies of all available documentation related to the system(s)'s potential to generate outputs that include hallucinations or fabrications. If no such documentation is available, please explain why not.					
1.8. Have any analyses of the system(s) been conducted, including but not limited to the following types of analyses? If so, please provide copies of such analyses or explain why such analyses are unavailable. - AI Red-teaming - Algorithmic impact assessments or other types of algorithm audits (including first-party, second-party or third-party algorithm audits)					
Part 2: Training Data Information *Please provide the following information about the data used to develop or deploy the model(s) or system(s) in question, including training data and any datasets used for finetuning, or explain why such information is not available. If information requested in this section is available in documentation you are providing from Section 1, you can indicate that in this section.*					
2.1. Provide summary information about the dataset(s), including: 1. A short description of the dataset(s), including the main sources from which the data was collected or assembled 2. The size of the dataset(s) 3. The rationale behind the dataset(s)'s creation 4. Who curated/created the dataset					
2.2. *[For datasets that include text and/or audio data]* Please describe the language distribution of the dataset(s) – all of the languages and language varieties represented in the dataset(s) as a percentage of the whole					
2.3. *[For datasets that include text and/or audio data]* Please provide information about the demographics					

ACLU– Generative AI Vendor Questionnaire

Question		Yes	No	N/A	Additional Comments or Clarifications	For ACLU Use Only
	of speakers/authors represented in the dataset (age, gender, race/ethnicity, socioeconomic status)					
2.4.	*[For datasets that include images, video and audio]* Please provide information about the demographics of content represented in the dataset (e.g., if the video or audio includes spoken language, the distribution of spoken language; if the video or audio includes depictions of humans, the distribution of demographics represented, etc.)					
2.5.	Do any current or past agreements for data sharing, data licensing, or other data acquisitions govern the dataset(s)'s collection and system development process? If so, please provide copies or explain why such information is unavailable.					
2.6.	Information about who annotated the dataset(s) (including those who developed the annotation guidelines and those who completed the annotation process), including about the working conditions of the data annotators.					
2.7.	Please describe the technical means by which the dataset(s) were collected (e.g., scraped from particular websites, assembled from publicly available datasets, etc.)					
2.8.	Please describe the process used to filter, document, or curate the dataset(s).					
2.9.	Were the owners or creator of content included in the dataset(s) explicitly notified about the use of their work, compensated for the use of their work, and given the opportunity to have their work removed from the dataset(s)? If so, please describe.					
2.10.	Does any content in the dataset(s) reflect social biases, stereotypes, discrimination or hateful content? If so, please describe.					
2.11.	Please provide any other documentation or resources you can provide about the dataset(s).					
Part 3: Privacy and Copyright						

ACLU– Generative AI Vendor Questionnaire

Question	Yes	No	N/A	Additional Comments or Clarifications	For ACLU Use Only
For relevant questions, please provide the following information about the data used to develop or deploy the model(s) or system(s) in question, including training data and any datasets used for finetuning, or explain why such information is not available. If information requested in this section is available in documentation you are providing from Section 1, you can indicate that in this section.					
3.1. Do the dataset(s) include copyrighted works or materials? If so, please describe.					
3.2. Were steps taken in the development process to prevent or address the inclusion of copyrighted works or materials in the dataset(s)? If so, please describe.					
3.3. How can we be sure we aren't engaging in copyright infringement when using the system(s)?					
3.4. Do the dataset(s) contain personal data linked or linkable to an individual? Do the dataset(s) contain sensitive data as defined by consumer privacy laws. If so, please describe.					
3.5. Please describe all remediation steps or other measures you have taken or would take upon identification of personally identifiable information in the dataset or system outputs.					
3.6. Have any analyses of potential training data leakage or any efforts to quantify data leakage been conducted? If so, please summarize and provide copies of such analyses or explain why they are not available.					
3.7. Do the system(s) or related artifacts include mitigations or measures to prevent training data leakage? If so, please describe.					
3.8. Have any analyses of potential PII leakage in user prompts or any efforts to quantify data leakage in user prompts been conducted? If so, please summarize and provide copies of such analyses or explain why they are not available.					
3.9. Do the system(s) or related artifacts include mitigations or measures to prevent data leakage in					

ACLU– Generative AI Vendor Questionnaire

Question	Yes	No	N/A	Additional Comments or Clarifications	For ACLU Use Only
user prompts? Do they contain specific safeguards against leaking prior user inputs or sensitive information? If so, please describe.					
3.10. Do you have a written policy or policies addressing the consumer privacy laws? If so, please provide copies of all such policies.					
3.11. Can people request data subject rights including opting out of specific data processing by your vendor, deletion from your training dataset and access to any personal data you have in your training datasets or models? If so, please describe.					
3.12. Can you provide contractual guarantees that our data won't be used to train the system(s) or stored and used for training of other system(s) without our specific consent? If so, please describe.					
3.13. As more states around the U.S. pass privacy laws, please describe how your approach to these issues is changing or updating.					
Part 4: Security					
4.1. Does the system include mechanisms for logging how users are interacting with the system, or can such mechanisms be added on to the system? If so, please describe.					
4.2. Are admin actions logged?					
4.3. Can logs be routed to third-party SEIMs or exported for audit?					
4.4. Is it possible to integrate systems for filtering user prompts before user prompt data is sent to the model/vendor's system(s)? If so, please describe.					
4.5. Do employees at the vendor have access to the prompts and responses submitted by users? If so, under what circumstances is this data accessed?					

ACLU– Generative AI Vendor Questionnaire

Question	Yes	No	N/A	Additional Comments or Clarifications	For ACLU Use Only
4.6. Has the AI system undergone red teaming, penetration testing, or adversarial testing specific to its generative AI features?					
Part 5: Bias and Fairness Evaluations					
5.1. Has the system(s) been evaluated with regard to bias and fairness issues and other social and ethical risks? If yes, please provide copies of all such evaluations.					
5.2. For any such evaluations, please describe who conducted the analysis (e.g., were they an independent entity without financial ties to the entity that built or deployed the system(s)?).					
5.3. Did the evaluation(s) include specific testing focused on the possibility for system outputs to include biased, discriminatory, or otherwise harmful content? If so, please describe.					
5.4. If the evaluations included testing focused on algorithmic bias, did the evaluation include a consideration of potential bias related to protected characteristics including but not limited to race/ethnicity, sex, age, disability, pregnancy status, and intersectional combinations of those characteristics? If so, please describe.					
5.5. For each evaluation, please describe the specific metrics and datasets used in the evaluation.					
5.6. For each evaluation, describe when the evaluation occurred relative to the system's design and deployment (e.g., the evaluation occurred during the design process and before deployment, after initial deployment, one year after deployment, etc.)					
5.7. For each evaluation, were any changes, updates, or decommissioning, or other decisions made to the system(s) or related artifacts as a result of the evaluation? If so, please describe.					

ACLU– Generative AI Vendor Questionnaire

Question	Yes	No	N/A	Additional Comments or Clarifications	For ACLU Use Only
5.8. Do you have processes that would allow us to request you conduct additional bias, fairness, or safety evaluations? If so, please describe such processes and the estimated turnaround time for completing such analyses.					
Part 6: Environmental Impacts *Given the significant amount of energy that generative AI systems use, we ask for information about the environmental footprint of the system(s) in question, the compute region where experiments to build the system(s) were run, and the steps taken to mitigate environmental impacts, especially on Indigenous communities.* *For resources on compute regions with low carbon intensity, please see CodeCarbon or Google Cloud Region Picker if you use a major cloud provider (Azure, Amazon Web Services, or Google Cloud Platform).* *For resources on AI's impact on Indigenous communities, please see the following: (a) these articles on the impact of data centers on Indigenous land, water, and power resources that exacerbate ongoing broadband access challenges; (b) these resources on consultations with Indigenous communities in data center development to honor tribal sovereignty (c) the Global Data Centers Database with information about the locations and characteristics of data centers globally; and (d) the Energy Access for U.S. Native Lands map with the locations and emissions of power plants, as well as the locations of protected Indigenous lands in the U.S.*					
6.1. In which compute or energy grid region were the experiments or other processes to build, modify, and deploy the system(s) run?					
6.2. Were steps taken to consult or mitigate the impact on Indigenous communities when developing and deploying the system(s)? This may include efforts to avoid reliance on data centers that negatively impact Indigenous communities. If so, please provide documentation, including any relevant information about the locations or characteristics of the data centers or energy grids used to build the system(s).					
6.3. Were steps taken to mitigate the environmental impacts of developing and deploying the system(s)? This may include efforts to minimize resource usage throughout the stages of the system's lifecycle including idle energy usage, model training, and model deployment. If so, please provide documentation related to measures taken to mitigate environmental impacts.					

ACLU– Generative AI Vendor Questionnaire

Question	Yes	No	N/A	Additional Comments or Clarifications	For ACLU Use Only
6.4. What is the carbon footprint in CO2 emissions from developing the system(s)? Please use this carbon footprint calculator if these numbers are not readily available or you are unsure how to conduct these measurements.					
Part 7: Other Deployment Details					
7.1. Please describe any safety procedures implemented with the system(s), including prompt and output filtering procedures, harmful content detection measures, etc.					
7.2. Is regular auditing or evaluation conducted in relation to the system(s)? If so, please describe the auditing/evaluation process, focus of the auditing/evaluation, and measures in place to take corrective or other action based on the auditing's findings. Please also provide copies of such audits or evaluations if available.					
7.3. Are the owners or creators of works used in the model's training data or finetuning data provided with compensation related to the model's use or deployment? If so, please describe.					