CFSA AI Values Alignment Report

Any CFSA employee seeking to use an Artificial Intelligence (AI) or Machine Learning (ML) tool or system must answer the questions in this form.

Case Name: Comprehensive Child Welfare Information System (CCWIS) Case Operations Resource Assistant (CORA) to be implemented in the new Stronger Together Against Abuse and Neglect DC (STAAND) system.

Section I. General Questions

1. Briefly describe your use case and its purpose.

In June 2025, CFSA will launch STAAND II, an improved version of its current case management system. Staff will be trained both on the new system and its built-in new AI tool, the Comprehensive Child Welfare Information System (CCWIS) Case Operations Resource Assistant (CORA). CORA is an AI-powered chatbot that provides immediate answers to staff questions about policies, procedures, and system functionality when staff are managing a case. The knowledge sources used by the chatbot are limited to documents relevant to process and policy that are currently separate from the STAAND case management system. The chatbot does not draw from confidential case files.

CORA will:

- Utilize natural language processing to interpret and respond to queries using approved documentation;
- Provide step-by-step guidance for common procedures and system operations;
- Provide guidance on documentation requirements for different case types; and
- Assist with form completion by providing relevant information and validation.

The goals of CORA are to:

- Reduce training time and calls to help-desk for staff;
- Increase consistency in policy and procedure implementation; and
- Provide more time for staff to focus on client interaction.

Consistency will be increased because currently staff need to wade through voluminous documentation and training tools to find answers to questions, which may cause variations in work product. Retrieving those answers from a chatbot trained to answer them accurately and uniformly should increase the consistency of case managers' work.

2. What kind of AI tool is being used, and how is it being used?

This is a chatbot model in Microsoft Copilot Studio, ChatGPT 4.0 Turbo. It will generate explanatory text that staff can follow and use as guidance for their work. It does not use machine learning or create new content. Its answers are limited to the approved knowledge base.

a. Indicate whether your AI tool will involve the synthetic generation of text, images, or video.

It will generate explanatory text that the employee will follow as guidance, not text to be used by the employee as part of case documentation. At this time, it will not generate images or video.

b. If yes, will the AI-generated material be made available to the public?

No.

3. What kind of data does the AI tool or system(s) need to function? Please provide:

a) Name of dataset(s).

CORA will use a data set in CFSA SharePoint pulled from these sources:

- CFSA website – cfsa.dc.gov;
- CFSA policy index – available here;
- CFSA SharePoint site containing STAAND training material and tip sheets

The data pulled from those sources will be validated for accuracy by the appropriate Deputy Director-appointed subject matter expert from the CFSA AI Steering Committee (for example, the Deputy Director of Policy, the General Counsel, etc). Once validated, the sources will be assembled on a SharePoint site and used as the knowledge source for CORA. As part of its AI monitoring responsibilities, the AI Steering Committee will review the content on an annual basis to ensure it is still accurate and to incorporate any policy or process updates.

b) Is the dataset open or private?

Open.

c) Which division and supervisor owns and maintains the dataset?

Policies are maintained by the Office of Policy, Planning and Performance (Christian Gineste) and the training materials are maintained by the Child Information System Administration (CISA) (Belinda Barton). OCTO/OPI maintains the website.

d) What is in the data set?

The data set contains information about CFSA operations, processes, and training, as well as a comprehensive set of policies that govern CFSA's operations and ensure compliance with DC and federal law. It does not include confidential case files.

e) How is the data set accessed?

The data set is accessed through a SharePoint site with knowledge sources validated by the appropriate subject matter expert from the CFSA AI Steering Committee.

f) Is the dataset currently in the DC Enterprise Data Inventory? If yes, indicate its current EDI classification (reference to OCTO Data Policy).

No.

g) Is the data sensitive or from a sensitive domain?

No.

h) Is the data set subject to an existing retention schedule?

Yes. Policies are reviewed on an annual basis.

i) Where will the data set reside? Will it remain in its current location after the AI tool is launched?

The information currently resides on the CFSA website and intranet. Once it is validated and placed on the SharePoint site, it will not change locations after the launch of CORA.

4. Which CFSA Division or Program will own this AI system after it is released?

CISA

5. Which individual employees within the Division are responsible for this system, and who is the lead?

Suresh Chandran, the IT Manager in charge of DevOps at CISA, will be the technology system lead. Ashley Wharton, ISO, will be the security lead. And Sarah Koreishi, the CFSA AI Officer, will be the overall AI program lead.

6. Who will use the system?

All CFSA staff will use the system.

7. How was the system tested, using what data?

CISA is currently testing CORA using the data sets described above.

8. How will users of the system be trained?

All staff will be trained on STAAND II (including its AI tools) by the CISA training team. The STAAND training will include a user guide for CORA. Starting at the end of March, there will be a separate training for all staff on the safe and responsible use of AI, using content from KnowBe4. For new staff, the training will be included in their onboarding training.

9. Is the use rights-impacting (e.g., workforce management, health and risk assessments, referral for services, eligibility for services, benefits determinations, making child placement decisions (adoption, family, or foster care))?

Yes, tangentially (see next answer). The system is designed to assist with ensuring that staff work product and decisions are made in compliance with CFSA policies. Therefore, the staff's decisions will be procedurally guided by the tool, although staff will remain substantively responsible for their decisions. And those decisions will be rights-impacting, including, for example, health and risk assessments, referral and eligibility for services, benefits determinations, and making child placement decisions.

> **a) If yes, specify how, including whether the AI system is only tangential to rights-impacting decisions or actually makes rights-impacting decisions.**
>
> The system is tangential to rights-impacting decisions, meaning that the staff will still be making the decisions. The system will not make any rights-impacting decisions. But it will provide the staff with the information and task process that the staff needs to make the decisions.

10. What statutory protections are implicated?

The following statutory protections are implicated:

- DC Official Code §§ 4-251.01 – 4-251.27 (Grandparent and Close Relatives Caregiver Programs);
- DC Official Code §§ 4-301 – 4-361 (Adoption Programs) ;
- DC Official Code §§ 4-1301.01 – 4-1371.14 (Child Abuse and Neglect);
- DC Official Code §§ 4-1401 – 4-1424 (Placement of Children in Family Homes);
- DC Official Code § 7-241 (Sharing Information with DC Government Agencies);
- DC Official Code § 7-1201.01 (Mental Health Records);
- DC Official Code §§ 7-2101 – 7-2108 (Youth Residential Facilities);
- DC Official Code § 16-311 (Adoption Records);
- CFSA HIPAA policies;

- Any policies in the CFSA Policy Index relevant to the above statutes;
- The Health Insurance Portability and Accountability Act (HIPAA); and
- The Federal Educational and Privacy Act (FERPA).

a) Is the use consistent with those statutes?

As long as the information generated by the CORA tool is accurate and the staff follow the requirements generated by the CORA tool, the use will be consistent with those statutes.

11. Does the tool affect client due process rights (CFSA decisions for which clients have a right to appeal)?

Yes, but the system is tangential to due process rights-impacting decisions, meaning, that the staff will still be making the decisions. The system will not make any due process rights-impacting decisions. But it will provide the staff with the information they need to make the decisions. As long as the information generated by the CORA tool is accurate and the staff follow the requirements generated by the CORA tool, the use will be consistent with due process requirements.

12. Is CFSA required to notify the US Department of Health and Human Services about this tool or system?

Yes, CFSA is required to notify the Administration of Children and Families (ACF) about any technology implementation used to carry out child welfare operations. CFSA has made the notification via an Advance Planning Document (APD) submission.

13. Is CFSA required to notify labor unions about this tool or system?

Yes. STAAND II and CORA are on the agenda of the Labor Management Partnership Meeting on May 22, 2025.

14. Is there consistent and continuous human review while the tool is in use?

Yes. While using CORA:

- All automated actions will require human verification;
- Implementation of audit logs will be required for all AI interactions; and
- Regular performance reviews and validation of AI responses will be required.

a) Can humans sufficiently validate (i.e., review and approve AI generated recommendations) relevant AI outputs before they are enacted?

No. But the content in the data source is validated by appropriate stakeholders prior to being added to the data source. CFSA policies, procedures, and tip sheets all go through the rigorous review and approval process described in response to question 3. While validation may not occur at the time of output, it does happen before the content is added as a data source and then after use as part of the quality assurance process described above.

b) If yes, will that review and validation be documented? Explain.

N/A

15. Can humans override outputs?

Yes.

a) If yes, how?

If the proposed steps did not make sense, the staff person can contact the helpdesk and/or the training team to determine how to proceed with a task.

Section II. DC AI Values Alignment (from the OCTO Handbook):

1. Clear Benefit to Residents

a. Who will benefit from the AI tool?

Staff will benefit from CORA because they are using a new system and new business processes that will save time on administrative work, and also provide more accurate information for them to rely on for decisions. This will ideally result in more accurate decisions and more consistency in case handling. It should also decrease processing time for cases.

Accuracy will be increased because currently staff need to wade through voluminous documentation and training tools to find answers to questions, which may cause variations in the accuracy of work product. Retrieving those answers from a chatbot trained to answer them accurately and uniformly should increase the accuracy of case managers' work.

District residents will benefit from CORA because accurate and consistent processes will result in improved services and potentially fewer challenges to decisions/litigation. If staff process tasks more quickly, it will result in shorter wait times for services.

b. How have you weighed this tool's benefit to residents against reasonable alternatives?

The alternative to the use of the CORA tool is the status quo. Currently, staff have to sort through voluminous online records and stacks of paper with tip sheets. They also have to take days of training and refresher training to learn the processes and procedures (this training is separate from the training on how to use STAAND II). These processes are time-consuming; the launch of STAAND II in general and the CORA tool in particular is designed to make these processes more efficient and user-friendly for staff.

c. If the vendor's marketing material specific to the AI tool is available, please attach a copy for review.

Microsoft Copilot Studio 2024 Wave 2 Release Notes are available here.

2. Safety and Equity

a. What risks of direct or indirect physical harm might flow from your proposed tool?

There is no risk of direct physical harm as a result of the use of CORA. There is a low tangential risk of harm if there is inaccurate data in the knowledge source, or if staff do not follow their training on the use of the AI tool. For example, if the information generated by the tool is inaccurate and an inaccurate decision is made that results in leaving a child in an abusive or neglectful home, then that is an example of a risk of physical harm.

Nonetheless, CORA's human oversight and quality control requirements, the rigorous process to validate the accuracy of the knowledge/data source before it is entered into the tool described above in response to question 3, and the collective responsibilities of staff to appropriately investigate and evaluate a multitude of factors when making rights-impacting decisions, make this risk very low. The advantages of using the tool far outweigh the risks.

When being trained on the system, employees will be advised of the risk of harm that may result if the human oversight and quality control systems are not followed.

b. What risks of exacerbating inequity might flow from your proposed tool?

The risk of the tool exacerbating inequity is low, as long as the information generated by the CORA tool is accurate/free of bias, and the staff follow the requirements generated by the CORA tool. Subject matter experts from the AI Steering Committee

will review the knowledge source for bias when validating it for use with CORA. The tool itself will not make rights-impacting decisions. The staff will continue to make the decisions, which means that the risk of bias will flow from the human decision-maker, which is no different from the status quo.

The tool actually may decrease the risk of inequity because one of its goals is to improve the consistency of decision-making in rights-impacting decisions. If decisions are more consistent, that may lead to a decrease in inequity. In addition, all staff will have equal access to the tool and there will be multi-language support for diverse population needs.

3. Accountability

a. How will you ensure responsibility for all government action flows clearly to an appropriate DC government official?

Suresh Chandran, the IT Manager in charge of DevOps at CISA, will be the technology system lead. Ashley Wharton, ISO, will be the security lead. And Sarah Koreishi, the CFSA AI Officer, will be the overall AI program lead.

b. How will you measure the performance of the AI tool throughout its lifecycle?

The system will be regularly monitored and evaluated by the AI Committee through the following quality control measures:

- Monthly accuracy reports;
- Performance metrics; and
- Structured feedback loop from staff. This will include real-time feedback from staff during use, including the ability to flag questionable answers as they appear.

The AI Steering Committee also will monitor performance metrics related to timelines for training, service referrals, abuse and neglect investigations, and child placement, as relevant to whether CORA is furthering CFSA's performance in these areas.

In addition, CISA will monitor the system using Microsoft Copilot Studio's key performance indicators available here.

4. Transparency

a. What is your plan for public engagement?

CFSA's communications team is preparing messaging for the public and external stakeholders about the launch of STAAND II. That messaging will describe the anticipated benefits of the new system for CFSA staff and clients. It will also include descriptions of CORA, its anticipated benefits to CFSA staff and clients, and the guardrails described in this document to ensure safe and responsible use of AI.

b. How do you plan to label public AI-generated material?

The CORA is an internal, administrative tool. The AI-generated content is not likely to be released to the public. Because it uses open data, release to the public would not necessarily be harmful.

When employees use CORA, AI assistants will be clearly identified as non-human in all interactions. AI-generated answers will include the source of the information with a link to the document cited.

c. How do you plan to disclose to residents when they are interacting with non-human agents?

Residents will not interact with the CORA tool.

d. Can CFSA clients opt out of the system or appeal decisions made by the system?

No. CORA is an internal administrative tool; it will not make decisions.

5. Sustainability

a. What steps have you taken to ensure this deployment is cost-sustainable over the long term?

The use of the CORA tool is currently included in CFSA's budget. It funds 25,000 queries per month. This tool will not increase the number of core technology stack planned licenses. If costs increase over the long term, CISA will have to adjust the budget, depending on the prioritization of needs.

b. What consideration have you given to the environmental impacts of the deployment?

Use of the CORA will have a positive environmental impact. Use of efficient cloud-based services minimizes energy consumption and reduces the use of paper for documentation.

Microsoft also is committed to securing 3.5 million carbon credits over 25 years, aiming to balance the soaring carbon emissions fueled by AI advancements and become carbon negative by 2030. More information is available here.

c. How will the deployment impact job quality for your existing workforce?

The goal of CORA is to improve job quality for staff. They will benefit because it will save time on administrative work, reduce repetitive tasks, and also provide more accurate information for them to rely on for decisions. This should ideally result in more accurate decisions, more consistency in case handling, and decreased processing time for cases. It also provides an opportunity for staff skill development in AI-assisted workflows. Many staff may not have experience with the responsible use of AI. The training on STAAND II and CORA will provide that guidance, and their use of CORA will provide them with experience using and evaluating an AI tool.

d. Have you considered whether the deployment will displace existing DC employees?

CORA cannot replace CFSA employees because humans will still be required to make decisions while using the tool. It is intended to support, not replace, the workforce.

e. If a vendor fails to meet contractual obligations, what are the alternative options that exist to ensure there is no loss of service?

Without the service, all of the job knowledge resources will still exist. Staff would have to revert to manually using those sources to guide their work. In the event of a contract violation, the contract provides for standard DC default remedies.

6. Privacy and Cybersecurity

a. How will the technical aspects of your deployment promote and protect privacy?

CORA uses open data, not private data, so the risk of a privacy violation is low. All staff use of private data will continue to operate under existing CFSA confidentiality processes and policies. Employees will be advised during training to not enter private information into the chatbot. This warning also will be added to the screen while the tool is in use.

Microsoft CoPilot also functions within the privacy and cybersecurity requirements of CFSA's existing contract with Microsoft. Other DC government CoPilot users will not be able to access CFSA data because STAAND is access-protected. Users are

required to have a log-in and be part of the CFSA secure access group in order to use STAAND.

b. Describe any privacy-enhancing technologies the AI tool will have to promote privacy.

See answer to 6.a.

c. How will the legal terms of your deployment promote privacy?

Microsoft CoPilot functions within the privacy and cybersecurity requirements of CFSA's existing contract with Microsoft.

d. How will you notify people whose privacy is impacted by your deployment? Explain any privacy notice you intend to implement as part of your proposed AI deployment.

CORA does not use private data as a knowledge source. If a staff member enters private data into the query and somehow that private data is released, CISA will conduct an incident response as described in its incident response policy.

7. Configuration

a. Explain the AI tool's logging options and monitoring capabilities

CORA is developed using Microsoft Copilot Studio, which is native to the Microsoft Office and Azure/Dynamics environments. Microsoft provides inherent capabilities for monitoring described here, and also logs and monitors the following data:

Outcomes and engagement: Knowing the end-result of a conversation helps identify where CORA is succeeding and where it needs improvement.
Knowledge source use: Seeing how often knowledge sources are used helps understand how well CORA is able to provide answers to user questions.
User feedback: Reviewing user feedback helps identify new user scenarios and issues and make improvements based directly on what users are asking for.

b. Explain how the logging capabilities of the AI tool support general troubleshooting.

When CORA fails during a task, it provides an error message. A list of error messages and actions are maintained at Microsoft's "understand error codes" site.

c. Explain in technical detail how those logging capabilities will support any potential security investigations.

The logging data described above will be available during incident investigations.

8. *Interface*

a. Will the AI tool integrate with any systems and applications?

Yes. CORA will be embedded with the CFSA STAAND application and will interface with STAAND. CORA will not have access to the case data in STAAND.

b. Are these integrations internal or external to DC's information technology environment?

They are internal to DC's Azure environment. The data is stored in the DC Cloud.

c. Identify what data sharing agreements, if any, you intend to enter into

CFSA has multiple data sharing agreements including with OSSE, DHCF, Court Services, DOH, and DHS.

9. *Public-facing AI Tool*

This section is Not Applicable because the CORA is not public-facing.

a. Identify whether any AI tool involved in your deployment will be public-facing.
b. Identify whether the public will interact with the AI tool directly, as with a chatbot or AI-powered dashboard.
c. Identify whether the public will obtain outputs directly from the AI tool, or if there will be a DC government human in the loop.

d. Explain any technical safeguards you will incorporate into the AI deployment, including hard numerical limits on the number or complexity of queries, and any guardrails like scrubbing, filtering, or pseudonymizing of outputs

10. ***Active Support***

a. Explain your plans for managing the AI deployment in the event that the provider ceases to provide adequate support

If Microsoft abandons Copilot Studio, CFSA will stop usage and return to use of the knowledge source websites and hardcopy guides. If technical support for deployment is inadequate, CFSA will continue using internal technical support to complete deployment.

b. Will you discontinue use of the relevant AI tools? Will you arrange for your agency or some third party to assume responsibility for active support throughout the remainder of the AI deployment's lifecycle?

CFSA will not discontinue use unless they are not supported in GCC. CFSA will use third-party support for the deployment.

c. Explain the relevant AI tool providers' current and future support obligations.

The District has a three-year contract with Microsoft for support. CFSA has a services contract with Microsoft Industry Solutions for technical support that has three option years remaining.

d. Explain your change control strategy concerning regular general releases and agile releases to address serious vulnerabilities or functional issues.

CFSA has an established IT Steering Committee that works through prioritization of business requests for technology change. Additionally, all AI requests will be evaluated and approved by the AI Committee at CFSA.

e. Explain the relevant tools' planned standard patching schedules and maintenance windows

Copilot Studio follows the cadence of PowerApps and Dynamics platform updates. CFSA has been managing regular releases and has a DevOps team in place to address issues.

11. ***Risk Management***

a. Explain how your mapping, measurement, management, and governance of these risks will account for risks arising outside DC's information technology environment, including, for example, those risks arising within a cloud service provider's environment.

CFSA will adhere to its internal risk management policies and procedures to manage risk. Its practices are aligned with industry standards and frameworks to account for AI-specific risks, including the NIST AI RMF. NIST provides a structured approach to manage risks related to AI systems using four key functions: MAP, MEASURE, MANAGE, GOVERN.

i. **MAP:** Identify risks associated with using Microsoft Copilot (GCC) (e.g, data breaches, AI bias, compliance violations, etc).
ii. **MEASURE**: CFSA will assess and quantify risks using risk assessment frameworks as guidance to measure how Copilot manages risk. As a FedRAMP compliant cloud-based AI tool, CFSA will leverage the results of the third-party risk assessments.
iii. **MANAGE**: CFSA will continue to work to implement controls, such as updating contractual language and ensuring key security and privacy features are enabled within Copilot Studio, such as audit logging and monitoring, access control, and encryption to ensure compliance with applicable laws and policies.
iv. **GOVERN**: CFSA will continue to develop its governance structure to effectively manage all AI deployments, including developing additional policies and guidance, as needed.

By applying the NIST AI RMF, CFSA can identify, assess, manage, and govern risks beyond its direct IT environment, ensuring that risks from Microsoft's cloud infrastructure are properly accounted for. Since Microsoft GCC is compliant with government security standards, these frameworks help validate and manage AI risks effectively.

The CORA underlying technology is Microsoft Copilot Studio tailored for U.S. Government customers through Government Community Cloud (GCC). Copilot Studio US Government plans are designed to adhere to federal cloud service requirements, including the Federal Risk and Authorization Management Program (FedRAMP) accreditation at a High

Impact level. In terms of data segregation and access, customer content is stored within the United States, physically and logically separated from non-government plans and access is restricted to screened Microsoft personnel. Additional information can be found on the Microsoft Copilot Studio for U.S. Government customers US Government customers - Microsoft Copilot Studio | Microsoft Learn

b. Explain in technical detail any disaster recovery rating associated with any tool or platform that is a part of your AI deployment.

CFSA has instituted disaster recovery in geographically diverse locations. In terms of Copilot, please see Microsoft commitments here.

c. Explain how you will document internal temporary acceptance of a given risk, and what remediation steps you will apply in the period following temporary acceptance.

Temporary internal acceptance of risk will be documented using a security waiver. During the period following temporary acceptance, the agency information security officer (ISO), in coordination with the CIO, will develop a mitigation plan that will describe the risk, personnel assigned to address the risk, risk response plan, severity/priority level, and a remediation timeline that will coincide with the approved security waiver.

d. Explain your proposed incident reporting and response protocols.

Please see CFSA incident response policy.

CFSA will adhere to internal incident response and reporting policies and procedures for incidents involving the use of AI at CFSA. This includes incidents that may have resulted in harm to an individual, diminished civil rights or civil liberties of an individual or group of individuals, unauthorized release of PII or other sensitive information, reports of obscene, hateful or inappropriate content, or a cybersecurity breach. Procedures for managing such incidents are appropriately coordinated among relevant officials and align with and will not supersede existing incident reporting requirements, such as those related to privacy and security incidents.

CFSA will establish an AI Incident Response Team to provide quick, effective, and orderly response to AI incidents. The AI IR Team's mission is to address, handle, and resolve allegations of any harm caused by an AI incident. CFSA will coordinate incident response activities with OCTO SOC and report AI-related incidents in accordance with applicable policies.