

Acceptable Use of Generative AI Policy

Policy Review Committee

Policy Review Committee Member	Signature
Travis Vowles Director, IT Advancement	 Travis Vowles (Jul 23, 2026 11:07:37 PDT)
Danny Kim Assistant Vice President People & Operations	 Danny Kim (Jul 23, 2026 15:05:47 PDT)
Brian Zumbano VP, Chief Development & Stewardship Officer, Chief of Staff	 Brian Zumbano (Jul 23, 2026 11:12:50 PDT)

Terms and Definitions

For all terms and definitions related to this policy and other SDF policies, please refer to the SDF IT Terms and Definitions Document located in the [IT Policies](#) on the All Staff SharePoint site.

1. Acceptable Use of Generative AI Usage Policy

1.1 Goal

The goal of this policy is to establish the responsible, ethical, and secure use of Generative AI (GenAI) technologies within San Diego Foundation (“SDF”). AI must be used in a manner that aligns with SDF’s mission, values, and commitment to community trust while mitigating risks related to data privacy, security, and fairness.

This policy aims to ensure the responsible deployment and use of Generative AI technologies in a way that enhances operational efficiency, supports community engagement, and maintains SDF’s reputation for integrity and ethical stewardship. AI should be leveraged as a tool to augment human capabilities, improve decision-making, and drive innovation while safeguarding against misuse, bias, or harm to stakeholders.

1.2 Risk

The use of Generative AI introduces potential risks, including but not limited to data privacy violations, biased or discriminatory outputs, misinformation, over-reliance on automated decision-making, and regulatory non-compliance. SDF acknowledges these risks and commits to mitigating them through robust security controls, transparent AI governance, ethical AI training, and continuous oversight. Failure to appropriately manage these risks could result in reputational damage, legal liabilities, and erosion of community trust.

As AI capabilities expand — including through agentic AI systems capable of taking autonomous actions on behalf of users — SDF recognizes that risks will continue to evolve. This policy is reviewed annually to ensure it remains current with the AI landscape and SDF’s operational needs.

1.3 Scope

This Policy applies to: All SDF Users (“Users”), regardless of their locations: Regular full-time as well as temporary part-time employees, interns, independent contractors provided with specific access to SDF Systems, approved vendors, as well as third-party employees under contract, who have access to or involvement with SDF's business processes, systems, and/or SDF Data including Confidential Information and Personal Information. SDF includes the parent organization, SDF, as well as all current and future subsidiaries, including the Policy and Innovation Center (PIC) and the San Diego Housing Fund (SDHF).

All SDF Data (“Data”): Any data SDF collects, creates, processes, stores, or uses relating to SDF members, employees, and/or business records. This data includes SDF's Confidential Information, proprietary information, financial information, trade secrets, internal documents, application code, software, as well as SDF member/employee Personal Information created and/or accessed/processed by SDF for business purposes.

2. Rules

This section outlines the essential rules governing Generative AI (“GenAI”) usage within the San Diego Foundation (SDF). All Users, contractors, vendors, and partners are required to follow these guidelines to safeguard SDF's mission, protect sensitive information, and uphold ethical and regulatory standards. Failure to comply may result in disciplinary action, legal implications, or other organizational measures.

2.1 Approved AI Tools

SDF maintains a list of Generative AI tools that have been vetted and approved by the IT Security team. Only approved tools may be used to process SDF data, including sensitive, confidential, or personally identifiable information (PII). Approved tools have undergone security and, where applicable, privacy review to ensure they meet SDF's data protection standards.

Tool	Description
Claude (Cowork, Code & claude.ai)	Anthropic's AI platform approved for SDF use. Includes Claude Cowork (desktop automation), Claude Code (AI-assisted coding; IT activation required), and the Claude web/desktop assistant. Use only with your SDF-provided account.
Microsoft Copilot	AI assistant embedded within SDF's Microsoft 365 environment (Teams, Outlook, Word, SharePoint, etc.). Use only with your SDF-issued M365 account.
ChatGPT	OpenAI's AI assistant, approved for SDF use with data privacy protections in place. Use only with your SDF-provided account — not a personal or free account.

For the most current list of approved AI tools, refer to the Approved Software list in the [IT Policies on the All Staff SharePoint site](#).

Supervisor approval required. Before being provisioned an approved AI tool, staff must obtain approval from their supervisor for their tool selection. This helps ensure teams remain aligned and, where business needs require, are using the same tool.

Switching between tools. Staff may use any of the approved tools above and are encouraged to select the one best suited to each task. Staff who wish to change their primary AI tool may request a switch no more than once every 90 days, and no more than two times per calendar year. Requests to switch tools may be submitted to HappyFoxIT@sdfoundation.org.

One license at a time. To meet budget requirements, staff may hold only one AI tool license at a time. Requests for additional licenses require supervisor approval and may require funding to come from the division requesting the additional license(s).

Selecting the appropriate model. Refer to the model guide for the appropriate tool and be sure to select the appropriate model to keep environmental considerations and costs in line and reduce overage fees. Any extra usage allocations fees must be approved by the staff member's manager or supervisor, in addition to IT, for budgetary reasons.

Personal AI Tools & Accounts Are Not Permitted for SDF Business

Personal AI accounts, free-tier AI services, and any AI tool not provisioned by SDF IT must not be used for SDF business purposes or with any SDF data. This includes personal accounts on platforms that also have SDF-approved versions (e.g., a personal ChatGPT or Claude account). Use of personal AI tools for SDF work — regardless of the data involved — is a policy violation. Approved tools are accessible through your SDF-provided accounts only.

2.2 Acceptable Use Guidelines

The following Do's and Don'ts provide clear direction on how to responsibly leverage GenAI tools and ensure alignment with SDF values, privacy obligations, and legal requirements. These guidelines establish specific actions that are permitted (**Do's**) and prohibited (**Don'ts**) to maintain compliance, security, and ethical AI use. While these practical rules outline immediate expectations for AI use, they are reinforced by the broader **Principles of Responsible & Beneficial AI** in Section 2.4.

DO:

- Use GenAI tools responsibly as an aid for efficiency and innovation, ensuring outputs align with SDF's mission and values.
- Verify that any AI-generated content used is accurate, ethical, unbiased, and does not violate intellectual property, privacy, or regulatory requirements.
- Clearly indicate when AI-generated content has been utilized in reports, formal communications, or donor interactions.
- Review AI-generated recommendations or outputs with human oversight before making decisions based on them.
- When working with sensitive, confidential, or PII data, use only IT-approved tools (Section 2.1). Where possible, provide context by **linking or referencing files stored in Microsoft 365 (SharePoint, OneDrive)** rather than uploading or pasting data directly into AI tools.

- Use non-approved AI tools (e.g., free consumer AI services) for general ideation, brainstorming, and creative thinking, provided no sensitive, confidential, or PII information is shared.
- Inform supervisors when GenAI tools are used to complete tasks that could impact organizational policies or external communications.

DO NOT:

- Use GenAI tools for making employment-related decisions, including recruitment, hiring, promotions, or disciplinary actions.
- Upload, input, or share sensitive, confidential, or personally identifiable information (PII) into any non-approved AI tool or consumer-tier AI product.
- Use AI tools as a sole decision-making authority for critical organizational decisions without human review and validation.
- Integrate GenAI tools with internal SDF systems or workflows without formal approval from the IT Director.
- Present AI-generated work as original human-created content without disclosure.
- Rely on AI-generated content without fact-checking or ensuring compliance with SDF policies and regulatory standards.
- Upload SDF documents directly to AI tools when the same information can be accessed by linking files from Microsoft 365.

2.3 Agentic AI Usage

Agentic AI refers to AI systems capable of taking autonomous or semi-autonomous actions on behalf of a user — such as editing and creating files, executing code, sending communications, interacting with external services, or browsing the web. Claude Cowork and Claude Code are examples of approved agentic AI tools at SDF.

Because agentic AI can take consequential, and sometimes irreversible, actions, additional oversight is required. The following guidelines apply to all agentic AI usage:

DO:

- Review and explicitly approve AI-proposed actions before they are executed, particularly when they involve sending communications, modifying files, publishing content, or interacting with external systems.
- Ensure agentic AI tasks operate only within approved SDF systems and platforms.
- Apply the same data handling standards to agentic AI interactions as to all other SDF data processing.
- Grant agentic AI tools only the minimum access necessary to complete a specific task (principle of least privilege).

DO NOT:

- Allow agentic AI to execute irreversible actions — such as sending emails, deleting files, or publishing content — without explicit human review and approval at each step.
- Configure agentic AI workflows to operate fully autonomously on SDF data without periodic human oversight checkpoints.

- Grant agentic AI tools broad or standing access to core SDF systems (e.g., Salesforce, SharePoint, Outlook) without prior IT approval.

Any automated workflows involving agentic AI that connect to SDF's core systems must be reviewed and approved by the **IT Director** in advance. Submit requests to HappyFoxIT@sdfoundation.org.

2.4 Principles of Responsible & Beneficial AI

SDF is committed to **safe, ethical, and transparent** use of Generative AI (“GenAI”) to support day-to-day operations, including drafting documents, summarizing information, enhancing productivity, and improving communication. The following principles guide how AI should be used within SDF to ensure it benefits Users while maintaining security, compliance, and trust.

The Principles of Responsible and Beneficial AI are intended to maximize the AI benefits for San Diego Foundation while minimizing the risk of damaging over 50 years of trust established by our organization.

1. **Privacy and Security.**

Protect personal and sensitive data by following robust security standards, maintaining compliance with relevant data protection regulations, and respecting the privacy of donors, beneficiaries, and stakeholders.

2. **Data Ethics.**

Commit to ethical data collection standards, including analysis and usage practices, ensuring that the data used is accurate, relevant, and collected with proper consent.

3. **Inclusiveness.**

Actively address biases and disparities in AI outputs and responses by developing a review process to monitor and evaluate work based on:

- o Non-discrimination and prevention of bias
- o Representative and high-quality data
- o Fairness
- o Equality
- o Inclusiveness in impact
- o Inclusiveness in design

4. **Accountability.**

Share accountability with AI applications that we develop, deploy, or utilize, ensuring that they align with our organization's mission, values, and principles of community.

5. **Transparency.**

Be transparent in the development, deployment, and utilization of AI technologies, providing, requiring, or requesting clear explanations of AI methodologies, results, reporting, measurement, and potential impacts on participants. In addition, we will provide adequate visibility of our AI outputs when autonomous AI has been utilized.

6. **Continuous Learning.**

Commit to the responsible use of tested and untested resources while staying informed about the latest developments in Responsible and Beneficial AI, incorporating best practices into our work.

7. **Collaboration.**

Actively engage with and learn from our nonprofit and community foundation peers, sharing experiences, challenges, and successes in Responsible and Beneficial AI.

8. **Legal Compliance.**

Be aware of, and abide by, applicable laws, regulations, and best practices concerning AI development and operations pertaining to fundraising, data protection, and AI systems. This includes monitoring emerging AI-specific regulations such as the EU AI Act and evolving US federal and state AI legislation.

9. **Social Impact.**

Strive to maximize the positive social impact of AI while minimizing any potential harm by focusing on the needs of beneficiaries and communities.

10. **Sustainability.**

Consider the long-term sustainability and environmental impact of AI technologies and advocate for sustainable AI practices within our organization.

11. **Honor Lived Experience.**

Prioritize honoring real lives and lived experience of staff, beneficiaries, and communities and strive to utilize AI as a tool to bolster and support people in their lives and jobs rather than replace them.

12. **Acknowledge Fear.**

Acknowledge that fear around technology is real and not unfounded. SDF will focus on keeping continuous conduits of communication open to express reservations and fears in a safe space.

3. Non-Compliance

All San Diego Foundation Users are required to adhere to this Policy. Failure to comply with this Policy may result in disciplinary action, up to and including termination of employment, termination of contract, and/or civil penalties and criminal sanctions, depending on the circumstances.

4. Exceptions

Where systems, procedures, or processes are not able to meet the requirements of this policy and an appropriate justification exists, an exception shall be raised for review and approval. Any exceptions to this Policy must be approved by San Diego Foundation's **Director of Information Technology** in advance. Exception requests may be submitted to HappyFoxIT@sdfoundation.org.

5. Revision History

Version	Date	Author	Changes
1.0	February 21, 2025	Kirstyn Archer (CyberGuards)	Original — Final

		Travis Vowles (SDF)	
2.0	July 21, 2026	Travis Vowles (SDF)	Annual Review: Added Approved AI Tools section (Claude, Microsoft Copilot, ChatGPT); added Agentic AI Usage section; updated data handling rules to permit sensitive/confidential/PII in approved tools with M365 file-linking preference; added ideation carve-out for non-approved tools; added personal AI tools prohibition callout; updated Policy Review Committee; tri-brand header; format alignment with AUP; fixed typo in Principle 12; updated '50 years' reference.

Acceptable_Use_of_Generative_AI_Policy_v2.0

Final Audit Report

2026-07-23

Created:	2026-07-23
By:	Travis Vowles (tvowles@sdfoundation.org)
Status:	Signed
Transaction ID:	CBJCHBCAABAAz-blItBDiUqGxJFuWdu1QotLX8YIQ1v4V

"Acceptable_Use_of_Generative_AI_Policy_v2.0" History

-  Document created by Travis Vowles (tvowles@sdfoundation.org)
2026-07-23 - 5:54:45 PM GMT
-  Document emailed to Travis Vowles (tvowles@sdfoundation.org) for signature
2026-07-23 - 5:54:50 PM GMT
-  Document emailed to Danny Kim (dkim@sdfoundation.org) for signature
2026-07-23 - 5:54:50 PM GMT
-  Document emailed to Brian Zumbano (brianz@sdfoundation.org) for signature
2026-07-23 - 5:54:50 PM GMT
-  Email viewed by Danny Kim (dkim@sdfoundation.org)
2026-07-23 - 5:55:35 PM GMT
-  Email viewed by Travis Vowles (tvowles@sdfoundation.org)
2026-07-23 - 5:59:58 PM GMT
-  Document e-signed by Travis Vowles (tvowles@sdfoundation.org)
Signature Date: 2026-07-23 - 6:07:37 PM GMT - Time Source: server- Signature captured from device with phone number XXXXXXX7059 -
Signature Appearance Selected: MOBILE_DRAW
-  Email viewed by Brian Zumbano (brianz@sdfoundation.org)
2026-07-23 - 6:12:42 PM GMT
-  Document e-signed by Brian Zumbano (brianz@sdfoundation.org)
Signature Date: 2026-07-23 - 6:12:50 PM GMT - Time Source: server - Signature Appearance Selected: DRAW
-  Document e-signed by Danny Kim (dkim@sdfoundation.org)
Signature Date: 2026-07-23 - 10:05:47 PM GMT - Time Source: server - Signature Appearance Selected: DRAW
-  Agreement completed.
2026-07-23 - 10:05:47 PM GMT