

FOR CEOs AND EDS OF MISSION-FIRST ORGANIZATIONS

# The NFP Leader's AI Starter Kit

The platforms, the tools inside them, and the 25 terms you need to know. Built for leaders running organizations under \$5M who don't have time for theory.

*The question isn't which AI to use. It's which tool inside that AI, and for what job.*

## PART ONE

### The Platform Map

Five platforms. Dozens of tools. Most leaders use one, barely. This shows you what each platform actually contains — and what to use when.

<b>Claude (Anthropic)</b>	Best for nuanced reasoning, tone control, long documents, and complex instructions. The default recommendation for leaders who want to think alongside AI.	
<b>CLAUDE.AI CHAT — START HERE</b> Your daily thinking partner. Drafting, strategy, synthesis, analysis. <i>Grant narratives, board memos, donor outreach, preparing for hard conversations.</i>	<b>PROJECTS — BUILD THIS</b> Persistent workspaces with custom instructions and memory. Your reusable AI infrastructure. <i>Building AI roles: Executive Coach, Interim CAO, Board Prep Partner.</i>	<b>COWORK — TEAM-READY</b> Desktop app for non-developers. Automates file and task management without code. <i>Hand AI capability to your ops or program staff without a learning curve.</i>
<b>CONNECTORS</b> Links Claude to Google Drive, Notion, Gmail, Slack, and more. <i>Let Claude read and work with documents already in your systems.</i>	<b>CLAUDE CODE — DEV ONLY</b> Command-line tool for agentic coding. For technical staff or consultants. <i>Custom integrations. Not for most leaders — but useful to know exists.</i>	<b>CLAUDE IN CHROME / EXCEL / POWERPOINT — BETA</b> Agents inside your browser and Office tools. Still maturing. <i>Experiment — not ready to rely on yet.</i>

<b>ChatGPT (OpenAI)</b>	The most widely known platform. Strong general capability. Most of your staff have already tried it here. A strong alternative. The frameworks are identical.	
<b>CHAT — START HERE</b> Browser and mobile interface. Strong general-purpose capability across drafting, research, and analysis. <i>Same daily use cases as Claude — whichever your team is more comfortable with.</i>	<b>PROJECTS — BUILD THIS</b> Persistent workspaces with memory. Direct equivalent to Claude Projects (launched 2024). <i>Same role-based agents — if your team is already on ChatGPT, build here.</i>	<b>CUSTOM GPTS</b> Pre-configured assistants you can build and share across your team or publicly. <i>Give staff a configured tool with guardrails — without setup on their end.</i>
<b>CANVAS</b> Collaborative writing environment. Edit documents alongside AI in a shared view. <i>Longer documents where you want to refine and iterate with AI visible.</i>	<b>IMAGE GENERATION (DALL·E)</b> Image creation built into the chat interface. No separate tool needed. <i>Event graphics, social content, quick illustrations for reports.</i>	<b>GPT STORE / PLUGINS</b> Marketplace of third-party integrations and pre-built assistants. <i>Search before building custom — someone may have built what you need.</i>
<b>Copilot (Microsoft)</b>	Strongest when your org lives in Microsoft 365. Less capable for standalone reasoning but unmatched for Word, Excel, Teams, and Outlook integration.	
<b>COPILOT IN M365 — THE VALUE</b> Embedded inside Word, Excel, Teams, Outlook, and SharePoint. Works with your actual documents. <i>Drafting in Word, summarizing meetings in Teams, cleaning data in Excel — without copy-paste.</i>	<b>COPILOT CHAT</b> Standalone chat interface. Capable, but not the reason to choose this platform. <i>Quick tasks when you are already in the Microsoft ecosystem.</i>	<b>COPILOT STUDIO — IT REQUIRED</b> Build custom AI agents at the organizational level. Requires admin access. <i>Org-wide deployments once you have proven value at the personal level.</i>
<b>Gemini (Google)</b>	Strongest when your org runs on Google Workspace. Less consistent for standalone reasoning, but strong if your team lives in Google Docs and Gmail.	
<b>GEMINI IN WORKSPACE — THE VALUE</b> Embedded in Docs, Gmail, Drive, and Slides. Works inside your documents without leaving Google. <i>Drafting in Docs, summarizing in Gmail, generating in Slides — where your team already works.</i>	<b>GEMS — BUILD THIS</b> Saved assistants with custom instructions. Closest equivalent to Claude Projects for Google-first orgs. <i>Role-based agents if your org runs on Google.</i>	<b>GEMINI CHAT</b> Browser and mobile interface. Capable generally, but Claude and ChatGPT outperform for complex reasoning. <i>Quick research tasks and Google-integrated queries.</i>

<b>Perplexity</b>	A search-first AI that cites its sources in real time. One job: current information with sources. Not a replacement for your primary platform.	
<p><b>PERPLEXITY SEARCH</b></p> <p>Real-time web search with AI synthesis and citations. Stronger than Google for complex research queries.</p> <p><i>Sector research, funder landscape, policy updates, fact-checking before you publish or present.</i></p>	<p><b>WHAT IT IS NOT</b></p> <p>No persistent workspaces, no document drafting, no strategic coaching. Not a reasoning partner.</p> <p><i>One job: research. Use it before every funder meeting, board conversation, or briefing.</i></p>	

## The Starter Stack for a Sub-\$5M Organization

<p><b>Primary AI: your thinking partner</b></p> <p>Claude Pro or ChatGPT Plus (\$20/mo). Pick one and go deep before you explore others.</p>	<p><b>Integration layer — inside your existing tools</b></p> <p>Copilot if you run Microsoft. Gemini if you run Google. Not a replacement — a layer on top.</p>
<p><b>Research tool</b></p> <p>Perplexity Pro (\$20/mo). Keep it bookmarked. One job, done well.</p>	<p><b>Total monthly cost</b></p> <p>\$40–60/month. Less than most of your other software subscriptions — and it does more.</p>

## PART TWO

## 25 Terms Worth Knowing

Not a technical dictionary. These are the 25 terms that come up in funder conversations, board meetings, and team discussions — in plain language, for leaders who run organizations, not AI departments.

<p><b>01</b> <b>Large Language Model (LLM)</b> The technology behind ChatGPT, Claude, and Gemini. Trained on enormous amounts of text to predict and generate language. Not a search engine — it generates responses, which means it can be wrong.</p>	<p><b>02</b> <b>Prompt</b> The instruction or question you give to an AI. The quality of your prompt determines the quality of your output. Writing prompts well is the most important skill to develop first.</p>
<p><b>03</b> <b>Prompt Engineering</b> The skill of writing prompts that reliably produce useful outputs. Not coding — more like learning to give very clear instructions to a highly capable new hire.</p>	<p><b>04</b> <b>Project (Claude) / Custom GPT</b> A saved AI workspace with custom instructions, uploaded context, and persistent memory. The difference between a one-off chat and a configured AI role.</p>
<p><b>05</b> <b>System Instructions</b> The configuration text you write to define how a Project behaves — its role, its constraints, its knowledge of your org. Better instructions = better performance every time.</p>	<p><b>06</b> <b>Context Window</b> How much information an AI can hold and reference in one session. Think of it as working memory. If a conversation gets too long, earlier content may fall out.</p>
<p><b>07</b> <b>Hallucination</b> When an AI states something false with apparent confidence. LLMs generate plausible-sounding text — they do not verify facts. Always check specific claims before using them externally.</p>	<p><b>08</b> <b>Model</b> The specific AI brain you're using — Claude Sonnet, GPT-4o, Gemini 1.5 Pro, etc. Your \$20/month subscription gives access to the strongest current model.</p>
<p><b>09</b> <b>Token</b> The unit AI uses to process text — roughly 3-4 characters. Affects speed, cost, and context window size. Explains why long conversations have limits.</p>	<p><b>10</b> <b>Memory</b> How an AI retains information across sessions. In Claude Projects, memory is persistent. Without Projects, each conversation starts from zero.</p>
<p><b>11</b> <b>AI Agent</b> An AI configured to take sequences of actions — not just answer questions, but complete multi-step tasks. Your Executive Coach Project is a simple agent.</p>	<p><b>12</b> <b>Automation</b> An AI-powered process that runs without your direct involvement. The difference between using AI and having AI work for you.</p>

<p><b>13</b> <b>Integration / Connector</b> A link between AI and your existing tools — Google Drive, Slack, your CRM, your email. Lets AI work with the systems your org already uses.</p>	<p><b>14</b> <b>API</b> Application Programming Interface — how software systems communicate. When someone says "connect this via API," they mean a technical link that requires a developer.</p>
<p><b>15</b> <b>Workflow</b> A defined sequence of tasks AI can execute end-to-end. Workflows move AI from tool to infrastructure.</p>	<p><b>16</b> <b>AI Governance</b> Your organization's policies for responsible, compliant AI use — what's approved, what's not, how data is handled, and who decides.</p>
<p><b>17</b> <b>Multimodal AI</b> AI that processes more than text — images, PDFs, spreadsheets, audio. You can upload a financial statement or board deck and ask Claude to analyze it.</p>	<p><b>18</b> <b>Knowledge Cutoff</b> The date an AI's training ended. It won't know about events after that date unless it has live web search enabled. Always verify time-sensitive information.</p>
<p><b>19</b> <b>RAG (Retrieval-Augmented Generation)</b> AI that searches your own documents to answer questions. When you upload your strategic plan to a Claude Project and ask it questions, that's a basic form of RAG.</p>	<p><b>20</b> <b>Fine-tuning</b> Training an AI model on your specific data to customize its behavior. More advanced and expensive than Projects. Not necessary for most sub-\$5M orgs yet.</p>
<p><b>21</b> <b>Foundation Model</b> The base AI system applications are built on — Claude, GPT-4, Gemini. When someone says "we built on GPT-4," they mean they used OpenAI's foundation model.</p>	<p><b>22</b> <b>Zero-shot Prompting</b> Asking AI to complete a task without providing examples. Works well for most everyday tasks. When output quality is consistently low, switch to few-shot prompting.</p>
<p><b>23</b> <b>Few-shot Prompting</b> Giving AI 2-3 examples of the output you want before the actual task. Dramatically improves quality for complex writing and tone-matching.</p>	<p><b>24</b> <b>Data Privacy / No Training</b> The setting that prevents AI platforms from using your inputs to train future models. On Claude Pro: opt out in Settings. On Claude for Teams: on by default.</p>
<p><b>25</b> <b>AI Adoption</b> The organizational process of integrating AI into everyday workflows. Most leaders are stuck between personal fluency and team adoption. That gap is what the full scaffold addresses.</p>	

## Ready for the full scaffold?

This starter kit shows you the terrain. The LeadAI Scaffold (\$49) gives you the sequence — what to build first, what to build next, and how to move your team without losing what you built. Includes three Project prompt templates, a governance policy starter, and a 30-day checklist.

[nfpleadai.com](https://nfpleadai.com)

[Get the Scaffold →](#)