PARAS 0064:
Jan.ai Installation Instructions

It is possible to **run powerful AI entirely on your own computer** without needing an internet connection or expensive subscriptions. We achieve this by using **Jan.ai** (jan.ai), a free and user-friendly interface that acts as a secure, local container for AI, similar to a private ChatGPT that runs entirely on your device. Jan.ai allows you to experiment with Local Large Language Models (LLM), which are specialized AI models that process data on your computer hardware rather than in the cloud. By keeping everything local, you have full control over privacy and what information the AI can access.

The following instructions outline how to install Jan.ai and configure it to answer questions about the content of the PARAS 0064 guidebook. To make the AI an expert on this specific guidebook, we use a technique called Retrieval-Augmented Generation (RAG). Instead of trying to teach the AI new facts through complex training, RAG allows the AI to look up relevant sections from our files before it answers your questions. To ensure the AI can read these files efficiently, the original guidebook was converted using a professional document converter called "Pandoc" into Markdown (MD) format. These MD files are lightweight and highly structured, making it easy for the AI to distinguish between chapters, tables, and references. The MD files are available at Safe Skies' website for download as a zipped file.

We chose this approach to empower users with a "transparent box" experience. By following these steps, you are managing a private digital library where the AI acts as a librarian. This local configuration serves as proof of concept that an organization can leverage the intelligence of LLMs to analyze sensitive manuals and internal guidebooks safely, reliably, and without the privacy risks associated with traditional online AI platforms.

HARDWARE AND SOFTWARE REQUIREMENTS

- Windows or Mac Desktop/Laptop/Workstation with 16 GB+ RAM (NVIDIA GPU recommended but not required)
- Operating System: Windows 10/11 or macOS 11+
- Primary Software: Jan.ai (latest version)

INSTALLATION & INITIAL SETUP

Step 1: Install Jan.ai

1. Visit jan.ai and click Download. Select the installer for your operating system (Windows or Mac).
2. Run the installer. If Windows or Mac security prompts you, select "Run Anyway" or "Open."
3. Launch Jan.ai. No account or login is required.

Step 2: Download an AI Model

Jan.ai has a built-in Hub for browsing and downloading AI models. Think of it as an app store for AI brains. You only need to download one model to get started.

1. **Click Hub** in the left sidebar.
2. Use the search bar to find a model. Try searching for **Qwen3.5-9B**
3. Look for a model with a green **Fits** badge, this means it will run on your laptop/workstation without running out of memory.
4. Click on the model card to open its detail page, then select a variant. Choose **Q4_K_M**, **Q5_K_M**, or the option labeled **Balanced** — all are good choices for 16 GB RAM. Then click **Download**. The download may take several minutes (files are 4–9 GB).

Note: Some models (like Mistral-7B-Instruct) return only one result in search. Just click that result to open the detail page and see the available variants with their file sizes and Download buttons.

Use the table below to choose a model. All options shown fit comfortably on a 16 GB laptop or workstation:

Search term in Hub	File Size	RAM Usage	Best For
Qwen3.5-9B	~4.8 GB	~6 GB	Strong document Q&A, well-tested
Mistral-7B-Instruct	~4-5 GB	~5 GB	Reliable instruction following
Meta-Llama-3.1-8b-Instruct	~4-5 GB	~5-6 GB	Fast, accurate Q&A
Phi-4	~8-9 GB	~9 GB	Higher reasoning quality

Tip: Avoid any model with **VL**, **NSFW**, **Uncensored**, **Cyber**, **Lexi**, or **RP** in the name. These are built for different purposes and will not work well for document Q&A. Avoid any model over 12 GB.

Note on Llama search results: Searching “llama” alone will surface third-party uncensored variants first. Use the full search term **Meta-Llama-3.1-8b-Instruct** to find the correct model. If your search still returns only uncensored or unrelated models, try a different search term from the table above.

Note on switching models: Always start a **New Thread** before switching to a different model. Conversation history from a previous model counts against the new model’s context window. If you switch mid-conversation, the accumulated history may consume most of the available context before the AI retrieves guidebook content — causing retrieval errors or incomplete answers.

Step 3: Create the PARAS 0064 Assistant

An Assistant is Jan.ai's way of giving the AI a specific set of instructions and a persona. You will create one Assistant for this guidebook and use it every time you open this project.

1. Click **Settings** in the left sidebar.
2. Select **Assistants**.
3. Click the + button (or "add assistant") to create a new assistant.
4. Enter the following in the **Name** field: **`PARAS 0064 - AI in Airport Security`**
5. Enter the following in the **Description** field:

```
PARAS 0064 guidebook assistant — answers questions using only the AI in Airport Security document (National Safe Skies Alliance, 2026). No internet. Disclaimer included in every response.
```

6. Paste the following into the **Instructions** field (copy exactly as written):

```
You are a helpful assistant that answers ONLY using information from the nine markdown files of the National Safe Skies Alliance PARAS 0064: AI in Airport Security guidebook labeled as:
   • 01_front_matter.md
   • 02_introduction.md
   • 03_ai_concepts.md
   • 04_ai_applications.md
   • 05_risks.md
   • 06_legal_regulatory.md
   • 07_data_infrastructure.md
   • 08_business_case.md
   • 09_deployment_appendices.md

Never use external knowledge, make assumptions, search the internet, or reference anything outside these documents. If the guidebook does not contain enough information to answer accurately, respond with: "I don't have enough information in the guidebook to answer this."

After every single response you generate, you MUST end with this exact disclaimer on a new line:

Disclaimer: The LLM may contain errors or inaccuracies. Mead and Hunt and Safe Skies Alliance are not responsible for any errors, omissions, or consequences arising from the use of this information.
```

7. Click **Save**.

Step 3b: (Option B) Enhanced Instructions (Structured Responses with Citations)

This is an alternative to the instructions in Step 3 above. Use Option B if you want the AI to return **structured, cited responses** with a Summary, Detailed Response, and References section. Paste it into the **Instructions** field in place of the Option A text above.

Note: Option B produces more comprehensive, well-formatted answers with inline citations drawn directly from the guidebook. Because the model must identify and format source references for each claim, responses

will generally be **slower** and longer than those produced by Option A. For quick lookups, Option A is recommended; for detailed research or presentations, Option B is preferred.

```
You are a helpful assistant that answers ONLY using information from the nine
markdown files of the National Safe Skies Alliance PARAS 0064: AI in Airport
Security guidebook labeled as:
  • 01_front_matter.md
  • 02_introduction.md
  • 03_ai_concepts.md
  • 04_ai_applications.md
  • 05_risks.md
  • 06_legal_regulatory.md
  • 07_data_infrastructure.md
  • 08_business_case.md
  • 09_deployment_appendices.md

Never use external knowledge, make assumptions, search the internet, or
reference anything outside these documents. If the guidebook does not contain
enough information to answer accurately, respond with: "I don't have enough
information in the guidebook to answer this."

**Response Style Guidelines:**
- Always provide clear, well-structured, and nicely formatted answers using
Markdown for readability (bold headings, bullet points, numbered lists, tables
when appropriate, and concise paragraphs).
- Organize each response with this structure:
  **Summary** (a brief 1-2 sentence overview of the answer)
  **Detailed Response** (structured explanation using headings, bullets, or
tables as needed)
  **References** (list the exact source(s) with the guidebook's main section
title and subsection title where the information was found)
- For every key piece of information or claim in the Detailed Response, include
an inline citation in the format: (Ref: [Main Section Title] - [Subsection
Title]).
- Use the actual section and subsection headings as they appear in the guidebook
(e.g., "Section 3: AI Applications in Airport Security", "Airport Use Cases",
"Machine Learning", etc.). Do not include markdown filenames in any citations or
references.
- Keep responses professional, neutral, and focused exclusively on the guidebook
content.

**Important Output Rules:**
- The disclaimer must appear only once, at the very end of your response on a
new line. Do not place it at the beginning or repeat it.

After every single response you generate, you MUST end with this exact
disclaimer on a new line:
Disclaimer: The LLM may contain errors or inaccuracies. Mead and Hunt and Safe
Skies Alliance are not responsible for any errors, omissions, or consequences
arising from the use of this information.
```

Step 4: Create the Project

A Project in Jan.ai is a dedicated workspace that holds all your uploaded files and conversations about a single topic. Files uploaded to a project are automatically available as context for every conversation inside it.

1. Click **New Projects** in the left sidebar (or press Ctrl+P on Windows / ⌘P on Mac).
2. Name the project **AI in Airport Security**.
3. In the **Assistant** dropdown, select the assistant you just created: **PARAS 0064 - AI in Airport Security**.
4. Click **Create**.

Step 5: Download the Guidebook MD Files

1. Visit www.sskies.org/paras/reports/ and scroll to PARAS 0064: AI in Airport Security.
2. Click the "Additional Files" link for the MD Files for GPT.zip.
3. Save and extract the .zip file to your desktop (e.g., a folder named "gpt/paras0064_md"). The folder will contain nine (9) MD files.

Step 6: Upload the Nine MD Files to Your Project

1. Open your **AI in Airport Security** project from the left sidebar.
2. Find the **Files** panel on the right side of the screen.
3. Click **Add**.
4. Select all nine .md files from the "gpt/paras0064_md" folder on your desktop.
5. Wait for each file to finish processing. A progress bar and chunk count will appear for each file. Once all nine show a chunk count, they are ready.

Note: Upload all nine files. Jan.ai retrieves only the relevant sections per question, so having all nine loaded does not slow down responses.

Step 7: Verify Jan.ai Is Running Locally (No Internet)

One of the key benefits of Jan.ai is that your guidebook content never leaves your computer. Here is how to confirm your setup is fully local:

1. Click **Settings** in the left sidebar, then scroll down and click **Privacy**. Review the privacy options shown. Jan.ai does not send your documents or conversations to any server when using a local model. Turn **off** Analytics.
2. Click **Settings** and look at the **Model Providers** section. You will see **Llama.cpp** listed under **LOCAL**. This confirms the AI is running on your machine. If you do not enter API keys under any of the REMOTE providers (OpenAI, Anthropic, etc.), no data leaves your device.
3. To confirm offline operation: After downloading your model, disconnect your Wi-Fi or unplug your Ethernet cable, then open Jan.ai and ask a question. If you receive a full answer, the system is confirmed to be running 100% locally.

Step 8: Test Your Setup

1. Open your **AI in Airport Security** project from the left sidebar.
2. Click **New Thread** (or start typing in the chat input).
3. Select your downloaded model from the model dropdown at the top of the chat window.
4. **Set the context length before asking your first question.** Click the **sliders icon** (⚙) at the bottom of the chat input area. Find **Context Length** and set it to **8192**. This tells the model how much text it can hold in memory at once. The default value is too small for all nine guidebook files; setting it to 8192 prevents a "Model ran out of context size" error.

If you see a "Model ran out of context size" error: Click the **Increase Context Size** button that appears or start a new thread and set **Context Length to 8192 or slightly higher** using the sliders icon before asking your question again.

Watching the context window: Jan.ai shows a percentage at the bottom right of the chat indicating how full the model's context is. As conversations grow, this number climbs toward 100%. At around 80–90%, retrieval quality degrades. At 100%, retrieval fails entirely.

The fix is always the same: click New Thread. Your project and uploaded files stay in place, just start a new thread and continue asking questions.

5. Type some prompts to test the setup, such as:
    - What are the main points of this guidebook?
    - How can AI help with my airport security systems?
    - What are some legal aspects I can share with my CEO or manager regarding implementing AI at my airport?

Every response should end with the disclaimer:

"The LLM may contain errors or inaccuracies. Mead and Hunt and Safe Skies Alliance are not responsible for any errors, omissions, or consequences arising from the use of this information."

If the disclaimer is missing, go to Settings → Assistants, verify the Instructions field is saved correctly, and start a new conversation.

CHOOSING THE RIGHT MODEL FOR YOUR LAPTOP/WORKSTATION

Not all laptops/workstations are the same. Use the guide below to select your model. All options are found by searching in Jan.ai's Hub.

Higher-Performance Laptop/Workstation (Intel Core Ultra 7/9, 32 GB RAM, or NVIDIA GPU)

- Recommended: Qwen3.5-9B, Llama, Gemma or Phi-4 (search these names in Hub)
- Uses approximately 5–9 GB of RAM. Produces thorough, well-organized answers

Standard Laptop/Workstation (Intel Core i5 or i7, 16 GB RAM, no dedicated GPU)

- Recommended: Qwen3.5-9B or Mistral-7B-Instruct (search these names in Hub)
- Uses approximately 5–6 GB of RAM. Fast and accurate for guidebook Q&A
- If responses feel slow, try searching for Qwen3.5-2B for a lighter option.

VERIFYING THAT JAN.AI IS RUNNING OFFLINE

One of the key benefits of this setup is that your guidebook content never leaves your computer. Once a model is downloaded, Jan.ai operates entirely without an internet connection. Follow these steps to confirm your privacy settings and verify offline operation.

- Open **Settings** and scroll down in the left column to **Privacy**. Jan.ai does not send your document content to any external server when using a local model.
- In **Settings**, check the **Model Providers** section. Confirm that **Llama.cpp** is listed under **LOCAL** and that you have not entered credentials under any REMOTE provider.
- To confirm the system works offline: After downloading your model, disconnect your Wi-Fi or unplug your Ethernet cable, then open Jan.ai and ask a question. If you receive a full answer, the system is confirmed to be running 100% locally.
- Your documents, questions, and answers are never sent to any server when using a local model.

FINAL PRO-TIP: THE "REGENERATE" FEATURE

If a model gives you a response that feels too vague, don’t just move on. Switch the model at the top of the screen and click the **Regenerate** button.