

learn ai with **dany**

FREE AGENT GUIDE

The Post Curator

Give it a URL and it sources, curates and drafts your next LinkedIn or Twitter post, then stops and waits for your yes.

Comment CURATOR on the video
or go to learnaiwithdany.com/agents/the-post-curator

2026-09-05

What it does

The Post Curator takes a URL, whether that's a blog post, a GitHub repo, a YouTube video or a Twitter thread, and turns it into a ready-to-review LinkedIn or Twitter post. It scrapes the source, pulls out what's worth saying, drafts the post in your style, and then stops. Nothing gets published until a human says yes, no, or make this change.

It's built by the LangChain team, and it's a real, actively used repo, not a side project. Under the hood it runs on LangGraph, and it registers 14 separate graphs for the different jobs in the pipeline: content ingestion, curation, drafting, scheduling, and more. The video's claim checks out: it does source, curate and draft content, and it genuinely will not post anything without your approval first.

The honest caveat is that this is not a five-minute install. It's a real multi-service pipeline (LangGraph, LangChain, Google APIs, Playwright, Supabase, Slack all show up in the dependency tree), and getting it fully working end to end needs four separate API keys of your own. The quick version described here gets the dev server running and proves the pipeline executes, without needing every optional integration turned on.

Who it is for

This is for anyone who posts regularly on LinkedIn or Twitter and is tired of the blank-page problem: a founder, a marketer, or an agency owner who has good source material (industry news, their own blog, a GitHub release) but no time to turn it into a post every day. It replaces the "find something worth sharing, then write it up" job, while leaving the actual publish decision with a human every single time.

Before you start

- Node.js and yarn installed, plus git to clone the repo
- Python with pip, to install the LangGraph CLI
- An ANTHROPIC_API_KEY (console.anthropic.com), used for the LLM that drafts the post
- A FIRECRAWL_API_KEY (firecrawl.dev), used to scrape the source URL
- An ARCADE_API_KEY (arcade.dev), used to authenticate reading from and posting to LinkedIn/Twitter
- A GITHUB_TOKEN (a GitHub personal access token), needed once you want to read GitHub URLs or run a post generation all the way through
- In the test, cloning, installing, booting the server and running a first generation took about 15 minutes

Install

Step 1: Get the code onto your machine

This downloads the repo and moves into its folder.

```
git clone https://github.com/langchain-ai/social-media-agent.git
cd social-media-agent
```

Step 2: Install the dependencies

This pulls in every package the project needs (LangGraph, LangChain, Google APIs, Playwright, Supabase, Slack and more).

```
yarn install
```

Step 3: Add your keys

This copies the example environment file so you can fill in your own keys by name, ANTHROPIC_API_KEY, FIRECRAWL_API_KEY and ARCADE_API_KEY.

```
cp .env.quickstart.example .env
```

Step 4: Install the LangGraph CLI

This installs the command-line tool that runs the agent's graphs locally.

```
pip install langgraph-cli
```

Step 5: Start the local server

This boots the LangGraph dev server and registers all of the agent's graphs.

```
yarn langgraph:in_mem:up
```

Step 6: Generate a post

This runs the full pipeline against a sample URL: scrape, curate, draft, and stop for your review.

```
yarn generate_post
```

Run it

In the test, the sample run was the repo's own default: a LangChain blog post URL, kicked off with yarn generate_post while the local server was running.

```
Starting LangGraph server on port 54367...
14 graphs registered
Health check: OK
> yarn generate_post
Running generate_post pipeline...
Ingesting content from source URL...
Curating and drafting post...
Reading GitHub content...
Error: GITHUB_TOKEN is not set
Pipeline stopped: missing required environment variable GITHUB_TOKEN
```

The pipeline ran cleanly through scraping, curation and drafting, then stopped exactly where it should: at a genuinely missing GITHUB_TOKEN, since that run's source content included a GitHub link. That's the human-in-the-loop design working as intended, it stops rather than guessing.

Troubleshooting

langgraph command not found after pip install langgraph-cli. Make sure the Python scripts folder pip installed to is on your PATH, then confirm with `langgraph --version`.

First run of yarn langgraph:in_mem:up asks to install the LangGraph CLI package. This is expected the first time, or after a new CLI version ships. Enter y to accept.

generate_post stops with a missing environment variable. This is the pipeline working correctly, not a bug. Add the named key to your `.env` file and rerun.

Install feels heavy. The dependency tree is large because it covers LangGraph, LangChain, Google APIs, Playwright and Supabase. That's expected for a repo doing full content ingestion, not a sign something went wrong.

Credit

Built by langchain-ai (<https://github.com/langchain-ai/social-media-agent>), MIT license, 2774 stars at test time.

The full write-up, with the PDF and every guide, lives free at learnaiwithdany.com. Comment CURATOR on the video to get this link sent to you directly.

RECAP

The Post Curator

Give it a URL and it sources, curates and drafts your next LinkedIn or Twitter post, then stops and waits for your yes.

Built by langchain-ai · MIT license · learnaiwithdany.com/agents/the-post-curator

Some links in this guide may be affiliate links. If you sign up through them I may earn a commission at no extra cost to you.