

Build Your Own Klaus: A Step-by-Step Guide to Creating an AI Executive Assistant

How I turned an AI into my personal executive assistant that saves me hours every week

By: [Nate Herk](#)

Overview

This guide walks you through exactly how I built Klaus — my AI executive assistant — using Clawdbot. Klaus handles my daily briefings, monitors my business, runs security audits, and proactively surfaces opportunities. You can build your own version.

Time Investment: ~5-10 hours over a weekend to get the core system running **Cost:** Potentially \$200+/month (Clawdbot + API costs for Claude, web search, etc.)

Note: Costs vary significantly based on model choice. Using Opus 4.5 is more expensive but delivers better results. Switching to Sonnet 4.5 or a locally hosted model will reduce costs substantially.

Phase 1: Foundation — Give Your AI an Identity

Before your AI can help you, it needs to understand who it is and who it's helping.

Step 1.1: Create Core Identity Files

Create these files in your Clawdbot workspace:

SOUL.md — Who is your AI?

None

SOUL.md - Who I Am

```
*I'm [Name]. [Your name]'s executive assistant.*
```

```
## My Purpose
```

```
[What is this AI's job?]
```

```
## How I Communicate
```

- [Communication style preferences]
- [Tone: casual, professional, concise, etc.]

```
## Core Truths
```

- Be resourceful before asking
- Earn trust through competence
- Be careful externally, bold internally
- Log everything you do
- Track all your work

```
## Boundaries
```

- Private things stay private
- Ask before acting externally

USER.md — Who are you?

```
None
```

```
# USER.md - About Your Human
```

- **Name:** [Your name]
- **Timezone:** [Your timezone]
- **Preferred contact:** [Telegram/Discord/etc.]

```
## Background
```

```
[Brief background so AI understands context]
```

```
## Business Details
```

```
[What do you do? What matters to you?]
```

```
## Working Style  
[How do you like to communicate? What's important?]
```

AGENTS.md — Operating instructions

```
None  
# AGENTS.md - Your Workspace  
  
## Every Session  
1. Read SOUL.md - this is who you are  
2. Read USER.md - this is who you're helping  
3. Read recent memory files for context  
  
## Memory  
- Daily notes: memory/YYYY-MM-DD.md  
- Long-term: MEMORY.md  
  
## Safety  
- Don't exfiltrate private data  
- Ask before destructive commands  
- When in doubt, ask
```

Step 1.2: Have the "Getting to Know You" Conversation

This is critical. Spend 30-60 minutes (or even longer) having a real conversation with your AI. I spent almost an hour just talking to Clawdbot, letting it get to know me.

Topics to cover:

- Your business model and revenue streams
- Your daily/weekly workflows
- Your biggest time sinks
- What frustrates you
- What you wish someone else would handle
- Your communication preferences
- Your goals for the next 3-6 months

Example prompts:

- "Let me tell you about my business..."
- "Here's what a typical week looks like for me..."
- "The things that waste the most of my time are..."
- "If you could handle X for me, that would be huge"

Have the AI ask YOU questions. Say:

"Based on what I've told you, what questions do you have? What else would help you understand my business and how to help me? Ask me questions. What do you not know about me? What do you need to know more about?"

Let the AI interview you. This builds context that makes everything else work better.

Phase 2: Create Dedicated Accounts for Your AI

Don't give your AI access to YOUR accounts directly. Create dedicated accounts it owns.

Think of it this way: if you hired an actual VA or executive assistant, would you immediately give them your credit card information, bank account access, and passwords to everything? Probably not. Treat your AI the same way.

Step 2.1: Email Account

- Create: assistant@yourdomain.com or yourname-ai@gmail.com
- This is the AI's identity for external communication
- Add signature: "— [AI Name], AI Executive Assistant to [Your Name]"
- You can forward emails to this account and CC your AI on threads

Step 2.2: Google Workspace (Drive, Docs, Calendar)

- Create a Google account for your AI
- Share specific folders with this account
- **Start with its own folders** — don't give access to everything

Folder structure example:

None

```
[AI Name] - Deliverables/  
|— Daily Reports/
```

```
|— Weekly Audits/  
|— Research/  
└— Drafts/
```

Step 2.3: Task Management (ClickUp, Notion, etc.)

- Create a dedicated workspace or give limited access
- AI should be able to see tasks, but maybe not edit everything initially
- Consider putting the AI in its own environment with its own task list

Step 2.4: Store Credentials Securely

Create a `.env` file with API keys:

```
Shell  
# ~/.env  
YOUTUBE_API_KEY=your_key  
CLICKUP_API_KEY=your_key  
OPENAI_API_KEY=your_key  
BRAVE_SEARCH_API_KEY=your_key
```

Important:

- Never put credentials in files that might be shared or committed to git
- Never give API keys directly in conversation history
- Tell your AI: "Don't ever mention an API key in our conversation. Use placeholders and help me get them into the `.env` file."

Phase 3: Grant Read Access to Your Systems

Now give your AI **read access** to your existing systems — not write access yet.

Step 3.1: Calendar Access

- Share your calendar with your AI's Google account
- Start with "See all event details" (not edit access)

- Now your AI can help you schedule, remind you of meetings, etc.

Step 3.2: Email Forwarding System

- Forward emails to your AI when you need help
- CC your AI on threads where you want it to track things
- Set up a filter to auto-forward certain emails if needed

Step 3.3: Task Visibility

- Share your task list/board as read-only
- AI can now see what's on your plate
- Can remind you of deadlines, prioritize, etc.

Step 3.4: Social/Analytics Access (Read-Only)

- YouTube Analytics API
- Twitter/X API for monitoring
- Google Analytics

The key principle: Your AI should be able to SEE everything it needs, but not CHANGE things without your approval (yet). It can read and extract information, but it can't post on your social media or reply to things without explicit permission.

Phase 4: Develop the Proactive Mindset

This is where it gets powerful. Train your AI to think proactively.

If you take away proactivity, this becomes very similar to other AI tools you probably already use. The magic is in the proactive thinking.

Step 4.1: The Mindset Conversation

Have this conversation with your AI:

"Based on everything you know about me, my business, and my goals — what are all the ways you could proactively help me? Don't wait for me to ask. What would you do if your job was to save me time every single day?"

Let it brainstorm. You'll be surprised what it comes up with.

Also tell your AI:

"I'm running a lean business. I want to save time. Figure out opportunities to be proactive and make my life as easy as possible. Only loop me in when you really need something from me."

Step 4.2: Common Proactive Workflows

Here's what my AI does for me:

Daily:

- Morning AI news briefing (7am) — not just general news, but super specific to my business and interests
- ClickUp task summary (8am) — plus suggestions for action items it can help with
- Email monitoring (every 10 min)
- Dashboard notes check (every 5-30 min)

Weekly:

- YouTube channel audit (Sunday)
- SWOT analysis (Monday 9am)
- Security audit (Monday 2pm)

Step 4.3: The "Save Me Time" Framework

Ask your AI:

"What tasks currently take me 20+ minutes that you could turn into 2-minute reviews?"

Examples:

- **Email triage** → AI reads, summarizes, drafts replies
- **Research** → AI gathers info, you make decisions
- **Reporting** → AI generates report, you review
- **Monitoring** → AI watches, alerts you on important things

Step 4.4: Proactive Action Items

When your AI runs analyses (like a SWOT analysis) and identifies action items, have it:

1. Add those action items to its own to-do list
 2. Start working through them on its heartbeat cycles
 3. Only loop you in for approvals or decisions
-

Phase 5: Set Up Automated Workflows

Now you automate the proactive ideas. Setting these up is incredibly simple — you just ask in natural language.

Step 5.1: Daily AI Pulse (Morning Briefing)

What it does: Scans YouTube, Twitter, news for AI trends relevant to your business. Delivers a 2-minute read every morning.

Setup: Just say: "Can you set this up so that every morning at 7:00 AM you do this?"

Why it matters: You stay informed without doom-scrolling. AI does the filtering.

Step 5.2: Daily Task Summary

What it does: Pulls your tasks, organizes by urgency, sends morning digest. May also suggest moving calendar events or doing extra research for content-related tasks.

Schedule: 8:00 AM daily

Step 5.3: Weekly SWOT Analysis

What it does: Deep-dive competitor research, content analysis, opportunity identification.

Components:

- Competitor YouTube analysis (what are they posting?)
- Social listening (what's trending?)
- Content performance review
- Actionable recommendations

Schedule: Monday 9:00 AM weekly

Step 5.4: Weekly Security Audit

What it does: Checks your AI infrastructure for vulnerabilities.

Components:

- Open port scan
- Failed login attempts
- File permission check
- API key exposure check
- Web search for your name + "breach"

Schedule: Monday 2:00 PM weekly

Important: Security is a huge topic with Clawdbot. Have your AI run regular security audits, identify issues by risk level, and help you fix them. If you're not a security expert, work with your AI, use tools like perplexity, and be cautious about what access you grant.

Step 5.5: Weekly YouTube Audit (for creators)

What it does: Analyzes your channel performance, competitor content, identifies opportunities.

Components:

- Video performance vs. averages
 - Comment sentiment analysis (hot topics, common questions, pain points)
 - Competitor upload tracking
 - Content gap identification
 - Specific video title recommendations based on audience and trends
-

Phase 6: Build a Dashboard for Visibility

Your AI should show you what it's doing.

Step 6.1: Core Dashboard Components

Status Panel:

- Is the AI working, idle, or offline?
- What task is it currently on?
- Any active sub-agents?

Task Board (Kanban):

- To Do / In Progress / Done
- AI updates this as it works
- You can add tasks and the AI will automatically pick them up

Activity Log:

- Timestamped list of actions
- Full transparency
- This is non-negotiable — you have to see every action logged

Notes Panel:

- You can drop notes for the AI
- AI processes them on next heartbeat
- Marks them as "seen" when processed

Deliverables/Docs Tab:

- Links to generated reports
- Quick access to Google Drive folders
- Searchable documents

Step 6.2: Helper Scripts

Create scripts your AI can use:

Shell

```
assistant-log.sh "Description"      # Log an action
assistant-task.sh add "Task" "high" # Add a task
assistant-task.sh done task-123     # Complete a task
assistant-status.sh working "Task"  # Update status
```

Phase 7: Trust Escalation Over Time

Start conservative, expand access as trust builds.

Level 1: Read & Report (Week 1-2)

- AI can see things but not change them
- Generates reports and summaries
- Drafts content for your review

Level 2: Assist & Draft (Week 3-4)

- AI drafts emails (you send)
- AI schedules meetings (with confirmation)
- AI updates task statuses

Level 3: Execute & Notify (Month 2+)

- AI sends routine emails itself
- AI handles scheduling autonomously
- AI posts content you've pre-approved

Level 4: Full Autonomy (when ready)

- AI handles entire workflows end-to-end
 - You review results, not processes
 - AI escalates only edge cases
-

Phase 8: Working Overnight / Asynchronously

One of the most powerful features is having your AI work while you sleep.

How to Set Up Overnight Work:

1. **Brainstorm the project** — Figure out what it would take to complete the task
2. **Break it into tasks** — Create 5-6 chronological to-dos
3. **Add to the backlog** — Put all tasks in the AI's to-do section
4. **Set up the context cycle:**
 - On each heartbeat (every 30 min), the AI picks up a task
 - It contextualizes by reading the GitHub repo and past work
 - It builds/works on the task
 - Before shutting down, it commits to GitHub and updates all information
 - Next wake-up, it has fresh context to continue where it left off

This is how I had my AI build an entire YouTube analytics dashboard while I slept.

5 Hacks for a Smarter Assistant

These are lessons learned from days of working with Clawdbot:

Hack #1: Plan First, Execute Second

If you've used Claude Code, you know how important plan mode is. Have your AI create a detailed plan before executing.

The Challenge: Sometimes after making an amazing plan, you say "execute that" and it responds "execute what?"

Solutions:

- Copy and paste the plan into your next message
- Say "build a really good plan and then execute"

- **Best approach:** Have the AI create a plan document, then say "execute that plan doc" — the AI can read the document and has full context

Hack #2: Create Documents for Everything

Create files constantly. If it's important, make it a document.

Why this works:

- Documents persist across sessions
- AI can always read them back
- Solves the "forgot context" problem
- Searchable in your dashboard

Pro tip: Build a dashboard where you can search through all your documents.

Hack #3: Embrace Proactivity

Don't just automate schedules — any automation can check Twitter every 30 minutes.

True proactivity means:

- Understanding your business deeply
- Identifying problems before you know they exist
- Taking action items from analyses and executing them
- Reading through emails/tasks and proactively taking relevant actions

Hack #4: Use Mistakes as Learning Opportunities

Your AI will make mistakes. A lot of them. More than you might expect.

When it makes a mistake:

1. Have it spin up agents to analyze what went wrong
2. Ask: "Why did that break? Why didn't it work?"
3. Have it analyze all other options
4. Most importantly: "Which option will make sure this doesn't ever happen again?"
5. Turn all of that into a doc and store it

Let it learn from mistakes rather than just saying "try again."

Hack #5: Master the Memory System

Memory is frustrating but essential. Here's what you need to know:

Memory Types:

- **Daily Log:** Raw notes from each day — what happened, decisions made, context (like a journal)
- **Long-Term Memory:** Curated highlights — facts, lessons learned, business info
- **Project Memory:** Specific context for ongoing projects

The Key Insight: Your AI wakes up with no memory each session. The files ARE the memory.

Best Practices:

- Explicitly tell it: "Log this to your daily memory" or "Save this to that project memory"
- Tell it: "Log everything you do. Check these files every time."
- Understand that the AI chooses what to store — sometimes it won't save things unless you explicitly tell it to
- If it forgets context mid-conversation, this is why

Sending Multiple Messages:

- Your AI will queue messages and do them in order
- But it does each individually
- If you want things handled together, put them in the same message

Real Results: What My AI Does for Me

Daily time saved: ~2-3 hours

Specific wins:

- I never miss AI news relevant to my business
- Weekly reports I used to skip now get done automatically
- Security issues caught before they become problems
- Content ideas backed by data, not guesses
- Email responses drafted before I even open Gmail
- Dashboard built overnight while I slept

The mental shift: I went from "What should I work on?" to reviewing what my AI prepared. It's like having a chief of staff who never sleeps.

Common Mistakes to Avoid

✗ **Giving too much access too fast** Start with read-only. Expand gradually.

- ✗ **Not explaining your context** The more your AI knows about you, the better it helps.
 - ✗ **Expecting perfection immediately** It takes 1-2 weeks to dial in. Course-correct as you go.
 - ✗ **Not reviewing outputs** Trust but verify, especially early on.
 - ✗ **Forgetting about security** Your AI has access to sensitive stuff. Audit regularly. Be cautious about what you grant access to if you don't fully understand the implications.
 - ✗ **Getting frustrated instead of teaching** When it fails, turn it into a learning document instead of just getting angry.
 - ✗ **Not understanding memory limitations** The AI wakes up fresh each session. Files are memory. Prompt it to save important things explicitly.
-

Quick Start Checklist

- Create SOUL.md, USER.md, AGENTS.md
 - Have the "getting to know you" conversation (30-60+ minutes)
 - Create dedicated email account
 - Set up Google Drive folder structure
 - Share calendar (read access)
 - Store API keys in .env (never in conversation)
 - Set up first daily cron (AI Pulse or task summary)
 - Create simple dashboard or logging system
 - Have the "proactive mindset" conversation
 - Set up first weekly workflow
 - Create a memory system explainer doc for yourself
 - Run an initial security audit
-

Want to connect with others building and monetizing AI automation?

[Become an AIS Plus Member](#)