

Revision 17 — February 22, 2026

So You Want To Set Up Your Own OpenClaw Server

A Complete Guide for Everyone (Yes, Even If You've Never Done This Before!)

This is a quick-reference guide with the minimum steps needed to create an OpenClaw server. For a comprehensive course with video walkthroughs and detailed explanations, check out the full class on Udemy: [Easy OpenClaw: Create An Employee Who Works for Free](#)

Table of Contents

1. **Part A: Prerequisites & Accounts** — What you need before starting
2. **Part B: AWS Account & Server Setup** — Creating your cloud computer (includes getting your API key)
3. **Part C: Telegram Setup** — Get this ready before installing OpenClaw
4. **Part D: Installing OpenClaw** — Getting your assistant running
5. **Part E: Your First Hour** — Configuring your assistant through conversation
6. **Part F: Optional API Keys & Services** — Adding extra capabilities
7. **Part G: Getting the Most Out of Your Assistant** — Tips and features
8. **Part H: What Else Can Your Assistant Do?** — Next steps to explore
9. **Conclusion**
10. **Really Technical Stuff** — Troubleshooting (you probably won't need this)
 - Real-Time Email Monitoring

Part A: Prerequisites & Accounts

What is OpenClaw?

Think of OpenClaw as your own personal AI assistant that lives in the cloud. Unlike ChatGPT where you open a website, or Siri that lives on your phone, OpenClaw is like having your own private assistant that you can talk to through Telegram (a messaging app), and it remembers everything about you, your projects, and your preferences.

The cloud just means “a computer somewhere else” - in this case, a computer that Amazon rents to you, which stays on 24/7 so your assistant is always available.

Why would you want this? Because OpenClaw can: - Remember conversations from weeks or months ago - Help you with projects over time - Connect to your email, calendar, and other services - Run tasks in the background (like checking your inbox) - Be customized to match YOUR personality and needs

It's like having a super-smart assistant who never sleeps, never forgets, and works exactly the way you want them to.

What You'll Need Before Starting

Let's be honest about what this takes:

Time: About 1-2 hours for your first setup. (Don't worry, you can take breaks!)

Money: - A credit card for AWS (Amazon's cloud service) - New accounts get free credits that cover the server for several months - After credits run out, about \$5-10/month for the server - API costs (the AI "brain") will be \$15-50/month depending on how much you chat

Accounts You'll Create: - Amazon Web Services (AWS) account - where your server lives - Telegram account - where you'll chat with your assistant - Anthropic account - this provides Claude, the AI brain

Technical Skill Required: Honestly? None. But you will need to: - Follow instructions carefully - Copy and paste commands exactly as written - Not panic when you see a black screen with white text (that's just the Terminal, and we'll explain it)

Cost Expectations (The Real Numbers)

Let's talk money because surprises are no fun.

Free tier period: - AWS EC2 server: \$0 (covered by AWS credits for new accounts) - AWS data transfer: \$0 (well within free tier) - Anthropic API (Claude): \$15-50/month depending on usage - Light use (few messages per day): ~\$15/month - Moderate use (20-30 messages/day): ~\$25-35/month - Heavy use (constant conversations): ~\$50/month - **Initial deposit: \$50 recommended** (ensures you don't hit rate limits during setup) - Optional APIs: - OpenAI (voice transcription): ~\$2-5/month if you send voice messages - Brave Search: \$0 (free tier covers most people) - Google AI (image generation): \$0 (free tier)

After credits run out: - AWS EC2 server: \$5-10/month - Everything else stays the same

Total realistic first year: \$180-600 depending on how much you chat

Total realistic ongoing: \$250-650/year

Compare that to ChatGPT Plus at \$240/year, but without any memory or customization!

Things You'll Save During This Process

Keep a note handy — you'll be collecting these items as you go:

- AWS account password
- Anthropic API key (starts with sk-ant-)
- SSH key file (.pem) — save this somewhere safe!
- Your Telegram User ID (a number)
- Your Telegram Bot Token (a long string)
- AWS backup credentials (Access Key ID and Secret Access Key)

Don't worry about understanding these yet — we'll explain each one as we go.

Part B: AWS Account & Server Setup

Step 1: Creating Your AWS Account

AWS (Amazon Web Services) is where we'll rent a tiny computer in the cloud. Think of it like renting a storage unit, except instead of storing boxes, you're renting a computer that stays on all the time.

Here's how to sign up:

1. Go to aws.amazon.com
2. Click the big "Create an AWS Account" button (usually in the top right)
3. Fill in your email address - use one you check regularly
4. Choose an account name - this can be anything like "MyOpenClaw" or your name
5. Click "Verify email address"
6. Check your email and enter the verification code
7. Create a strong password (write it down somewhere safe!)
8. Select **Free** here.
9. Fill in your personal information:
 - Full name
 - Phone number
 - Address

- All the usual stuff
10. Payment information: Enter your credit card
 - Yes, they need a card even for the free tier
 - They'll charge \$1 temporarily to verify it's real (then refund it)
 - The free tier really is free, and we'll set up billing alerts so you never get surprise charges.
 11. Identity verification: They'll text or call you with a code
 - Enter the code they give you
 12. And you're done! You should see a screen saying that Amazon is setting up your account.

⚠ Wait a few minutes. Sometimes AWS takes a little while to activate new accounts, so we're going to go set up your API key while that's being taken care of.

Step 2: Getting Your Anthropic API Key

Your cloud server needs someone smart on the other end to have a conversation with. That's where Anthropic comes in. They make Claude, the AI that will power your assistant's intelligence and personality.

Think of an API key like a special password that lets your assistant talk to Claude's brain.

1. Go to **console.anthropic.com**
2. Enter your email address and click "Continue with email"
3. Check your email for a sign-in link from Claude Console and click it
4. Enter your full name, choose a nickname, agree to the terms, and click Continue
5. Choose **Individual**
6. You're in! This is the Claude Console — your home base for managing your API key and monitoring usage.

Now let's get your API key:

1. Click **API Keys** in the left menu (under "Manage")
2. Click **Create Key**
3. Give it a name like "OpenClaw" so you remember what it's for
4. Leave the workspace on Default and click **Add**
5. You'll see your API key — a long string starting with sk-ant -
6. **Click "Copy Key" immediately** and paste it into your notes
7. You won't be able to see this key again after closing this window!

Set up billing so your assistant can use the API:

1. Click **Billing** in the left menu
2. Click **Add credit card** and enter your payment information
3. Once your card is added, click **Edit** next to the auto-reload warning
4. Turn on **Auto-reload** and set whatever amounts you're comfortable with — this way your credits automatically refill when they get low
5. To set a spending limit, click **Limits** and scroll down — you can set email notifications at whatever amount you choose

Important: Start with at least **\$50 in credits** and set auto-reload to trigger at **\$20**. This isn't just about having enough money — Anthropic's rate limits are directly tied to your credit balance. A low balance means you'll hit "cooldown" warnings where your assistant temporarily can't respond, which is frustrating during setup when you're chatting frequently. A \$50 balance puts you in a comfortable rate limit tier where you won't run into this. Typical monthly usage is \$15-50 depending on how much you chat, so \$50 lasts most people at least a month.

✓ **Checkpoint:** You should now have your Anthropic API key saved (starts with sk-ant-). Keep it handy — we'll need it when we install OpenClaw in Part D.

Step 3: Launching Your Server (EC2 Instance)

Okay, here's where we actually create your cloud computer! AWS calls these "EC2 instances" (Elastic Compute Cloud), but just think of it as "renting a computer that stays on all the time."

Finding the right place:

1. In the AWS Console, click on the search bar at the top and type "EC2".
2. Click on the tiny letters **EC2** next to the icon — not the icon itself, just the text.
3. Click on **Dashboard** in the left sidebar.
4. Click the orange **Launch Instance** button.

Configuring your instance:

Name: Give it a name like "My OpenClaw Server".

Application and OS Images (the computer's operating system): - Choose **Ubuntu** in the Quick Start section. - (It usually defaults to the latest LTS version, which is perfect.)

Instance type: - Select **t3.small** from the dropdown. - Make sure it shows "Free tier eligible".

Key pair (login): - This is like the only key to your server — if you lose it, you won't be able to connect. - Click **Create new key pair**. - Name: my-

openclaw-key. - Type: **RSA**. - Format: **.pem**. - Click **Create key pair**. - It will automatically download a .pem file. **Hold onto that file!**

Network settings: - The defaults should already be correct (“Create security group” selected, SSH allowed from Anywhere). - You can click **Edit** to verify: Auto-assign public IP should be Enable, and SSH allowed on port 22.

Storage: - Change it from 8 to **20 GiB**. OpenClaw needs room to grow. - Then click the orange **Launch instance** button at the bottom right.

Success! You’ll see a green “Success” banner. Scroll down and click **View all instances**.

Your instance will appear in the list. Wait for the **Instance state** to show “Running” with a green check mark. That means your server is booting up!

Click on your instance name to see its details. Find the **Public IPv4 address** and copy it or write it down. You’ll need this address in the next lecture to connect to your server.

Step 4: Understanding SSH

SSH stands for “Secure Shell,” but forget the technical term. Think of it as a secure phone line to your server. When you connect, you’re “talking” to your server via text commands.

Why? Because your server doesn’t have a screen. Is it safe? Very. That .pem key file you downloaded is like a special ID badge. Nobody can connect without it.

Step 5: Connecting to Your Server

Open Your Terminal

- **Mac:** Open **Terminal** (Applications → Utilities, or Cmd+Space “Terminal”).
- **Windows:** Open **PowerShell** from the Start menu.

This is just plain text — no mouse. But it’s incredibly powerful.

Lock Down Your Key File (Mac Only)

Before connecting, we need to lock down that key file so only you can read it. You’ll need to know where you saved it (usually Downloads) and what you named it (e.g., my-openclaw-key.pem).

Mac users, type this command (replace with your actual file name):

```
chmod 400 ~/Downloads/my-openclaw-key.pem
```

(Windows users can skip this step.)

Connect!

Now type the SSH command. Replace YOUR-IP with your server's Public IPv4 address, and use your key file's name:

```
ssh -i ~/Downloads/my-openclaw-key.pem ubuntu@YOUR-IP
```

For example: `ssh -i ~/Downloads/my-openclaw-key.pem ubuntu@18.234.56.78`

Press Enter.

The "Are you sure?" Question

The first time you connect, you'll see a message about "authenticity of host" and a fingerprint. **This is totally normal!** It's just asking "Is this really your server?"

Type yes and press Enter.

You're in!

You'll see a welcome message from Ubuntu and a command prompt (ubuntu@...). This means you're typing commands directly on your cloud server!

Don't close this window — we'll use it in the next section to install OpenClaw.

Step 6: Create Backup Credentials (IAM)

Before we leave the AWS Console, let's create a "key to the supply closet" so your assistant can automatically back up your server later.

1. In the AWS Console search bar, type **IAM** and click the text result.
2. Click **Users** in the left sidebar, then click the orange **Create user** button.
3. Name the user MyAssistant. Leave the "Provide user access to the AWS Management Console" checkbox **unchecked**. Click **Next**.
4. Under Permissions options, click **Attach policies directly**.
5. In the search box, type AmazonEC2F. Select the checkbox for **AmazonEC2FullAccess**. Click **Next**.
6. Click **Create user**.

Now to get the keys:

1. Click on your new user (MyAssistant) in the list to open its details.
2. Click the **Create access key** link (usually in the Summary section on the right, or under Security credentials).
3. Select **Application running outside AWS**. Click **Next**.

4. Skip the description tag value and click **Create access key**.
5. **Copy both keys** — the **Access Key** and the **Secret Access Key**.
 - Save them in your secure notes immediately.
 - You can also click **Download .csv file**.
 - **Important:** This is the *only* time you will see the Secret Access Key.

You're done with the AWS Console! Keep your terminal open for the next part.

Part C: Telegram Setup (Do This First!)

The easiest way to communicate with your assistant is the same way you communicate with everyone else — through a messaging app. OpenClaw can connect to WhatsApp, iMessage, and other systems, but I recommend using Telegram. It's a messaging app like WhatsApp or iMessage, but it has special features like "bots" (automated accounts) that make it perfect for OpenClaw. It works solidly, it's free, and it keeps your assistant conversations separate from all your other texting traffic.

Before we install OpenClaw, let's get everything ready on Telegram. This way, when we run the setup wizard, you'll have all the information it needs.

Step 1: Get Telegram

If you don't have Telegram yet:

1. Go to telegram.org
2. Download it for your phone (iOS/Android) or computer (Mac)
3. Install and open it
4. Sign up with your phone number
5. Enter the verification code you receive via SMS

Step 2: Get Your Telegram User ID

Your Telegram ID is a unique number that identifies you. You won't need it during installation, but it's handy for troubleshooting and for advanced configuration later.

1. In Telegram, search for: **@userinfobot**
2. Click on it and press "Start"
3. It will immediately reply with your information, including your **Id** (a number like 123456789)
4. **Write down this number!**

That's it! One message, instant answer.

Step 3: Create Your Bot

Now let's create the bot you'll talk to:

1. In Telegram, search for: **@BotFather**
2. Click on the bot called "BotFather" with a blue verified checkmark
3. Click "Start" at the bottom
4. Type: /newbot
5. BotFather asks for a **display name**:
 - Type the name you'll call your assistant. "Claw" puns are best. I use "ClawBot" but "Clawdette" would be a good one!
 - This is what you'll see in your chat list
6. BotFather asks for a **username**:
 - This must end in "bot"
 - Must be unique across all of Telegram
 - Try: YourNameAssistantBot or YourName_OpenClaw_Bot
 - If it's taken, try adding numbers
7. Success! BotFather will reply with:
 - A message saying "Done!"
 - A link to your bot
 - A **token** that looks like: 123456789:ABCdefGHIjklMNOpqrsTUVwxyz
8. **SAVE THAT TOKEN!**
 - Copy it and paste it into a note
 - You'll need it in Part D
 - Don't share it with anyone - it's like a password

✅ **Checkpoint:** You should now have TWO things written down: - Your Telegram User ID (from @userinfobot) - Your Bot Token (from @BotFather)

Got both? Great! Now we can install OpenClaw.

Part D: Installing OpenClaw

This is the moment we've been building toward. We're going to install OpenClaw on your server. The good news? The OpenClaw team has made this incredibly simple — it's literally one command.

Step 1: The One-Line Install

You should still have your terminal open and connected to your server from Part B. If so, you'll see that `ubuntu@ip-something:~$` prompt and you're ready to go.

If you closed your terminal, no worries — just open it back up. On Mac, open Terminal. On Windows, open PowerShell. Then type the same SSH command you used before:

```
ssh -i ~/Downloads/your-key-name.pem ubuntu@your-server-ip
```

Replace “your-key-name” with whatever you named your key, and “your-server-ip” with the IP address you copied from the AWS console. You can find it on the Instances page if you need to look it up again.

Once you see that `ubuntu@` prompt, you're connected and ready.

Here's all you need to type. One command installs everything:

```
curl -fsSL https://openclaw.ai/install.sh | bash
```

That's it. This single command downloads the OpenClaw installer, which then installs everything your server needs — Node.js, Git, and OpenClaw itself.

Go ahead and type that command, or copy and paste it. Then press Enter.

You'll see text scrolling by as it works. First, it installs Node.js — that's the programming language OpenClaw is written in. You don't need to know anything about it; the installer handles everything.

Next, it downloads and installs OpenClaw itself. This part may take a few minutes. You might see some warnings in yellow text — that's normal. As long as the installer keeps going, everything is fine.

When the installation finishes, you'll see “OpenClaw installed successfully” and the setup wizard will start automatically.

Step 2: The Setup Wizard

Follow the prompts using your arrow keys and Enter:

1. **Security Warning:** Read it, then select **Yes**.
2. **Onboarding Mode:** Select **QuickStart**.
3. **Model Provider:** Select **Anthropic**.
4. **Authorization Method:** Select **Anthropic API key**.
5. **Paste your API Key:** Paste the Anthropic key you saved in Part B (starts with `sk-ant-`) and press Enter.
6. **Default Model:** Press Enter to keep the default.

7. **Messaging Channel:** Select **Telegram**.
8. **Bot Token:** Paste the token you got from @BotFather and press Enter.
9. **Skills:** Select **Yes**. Use arrow keys to browse and space bar to select any you want (or just press Enter to continue).
10. **Hatch:** Press Enter to accept the default.

Your assistant is now configured! The wizard automatically starts your assistant as a background service, so it's already running. You should see a line that says something like "🤖 Telegram bot connected as @YourBotName". There it is — your assistant is online!

Step 3: Your First Conversation

Open Telegram on your phone or computer. Search for your bot's username — the one ending in "bot" that you created with @BotFather. Click on your bot, then click "Start" at the bottom.

You'll see a message with a **pairing code**. This is a security feature — OpenClaw needs to verify that you're the owner of this bot before it'll talk to you. Note the code in the message.

Now **switch back to your terminal** where OpenClaw is running. Type this command into the terminal:

```
openclaw pairing approve telegram KVDS4GDB
```

replacing KVDS4GDB with your own code. Press Enter. Your assistant will identify it as a command and ask if you want it to execute it. Say yes. You'll see "Telegram sender approved" with a little celebration! 🎉

Now switch back to Telegram and type "Hello". And there it is — your assistant responds! It's excited to meet you and starts asking what you'd like to call each other.

Congratulations! You just had your first conversation with your own personal AI assistant!

Try asking it a few things:

- "What can you do?"
- "Tell me about yourself"
- "What's the weather like?" (if you added web search)
- "Help me understand how you work"

See how it responds naturally? That's Claude's intelligence combined with OpenClaw's capabilities. And it's running on YOUR server, under YOUR control.

Good news: The setup wizard already configured your assistant to run as a background service. That means:

- It starts automatically when your server boots
- It restarts itself if it ever crashes
- It runs in the background — you can close your terminal window

Go ahead and close your terminal. Your assistant is online 24/7 in Telegram! If your server ever reboots (which is rare), just wait a minute or two and your assistant will be back online.

Part E: Your First Hour

Now for the fun part — making this assistant actually yours!

Your assistant uses text files behind the scenes to remember who you are and how to behave. But you don't need to edit any files yourself. Just talk to it. Work through these messages one at a time — copy, paste, and send each one via Telegram. Each builds on the last.

Most of these take less than a minute. The whole list takes about 30 minutes, and you'll have a well-configured, secure, personalized assistant when you're done.

1. Give It a Name and Personality

Your assistant starts generic. Fix that first.

“Let's set up your personality. Your name is [pick a name you like]. Write a SOUL.md file that describes who you are — your name, your vibe, how you communicate. Be direct and helpful, not corporate or sycophantic. Have opinions. Be concise when the situation calls for it and thorough when it matters.”

What this does: Creates SOUL.md, which shapes how your assistant talks and behaves across every conversation. It persists through restarts.

2. Tell It About You

The more context your assistant has, the better it can help.

“Let's create my USER.md profile. My name is [your name]. I live in [city/timezone]. Here's what I do for work: [brief description]. My communication style preference: [e.g., 'be direct, skip the fluff' or 'explain things thoroughly']. I mainly use you via Telegram on my [phone/computer/both].”

What this does: Creates USER.md so your assistant knows your timezone, work context, and preferences. Add more details anytime by saying “Update USER.md with...”

3. Set Your First REMEMBER Rules

Whenever you say “REMEMBER” (in all caps) followed by information, your assistant immediately saves it to long-term memory. This works in every conversation, forever. It’s one of the most powerful features you have. Here are three good ones to start with:

“REMEMBER: Always present times in my local timezone, not UTC.”

“REMEMBER: When I ask you to do something and you’re unsure, ask me a clarifying question rather than guessing.”

“REMEMBER: If you’re not sure about something, say so. Don’t make things up.”

4. Secure Your Server

These three messages harden your setup against attacks.

“Install fail2ban and configure it to ban IP addresses after 3 failed SSH login attempts for 24 hours.”

“Audit your own security setup and tell me if anything needs attention. Check SSH configuration, open ports, system updates, and anything else that might be a vulnerability.”

“REMEMBER: Run a security self-audit once a month and let me know if anything needs attention.”

What these do: Fail2ban blocks brute-force attacks automatically. The audit checks your current security posture. The REMEMBER ensures ongoing vigilance.

5. Set Up Monitoring and Alerts

Your assistant can watch itself and only bother you when something’s actually wrong.

“Set up monitoring to alert me if: you stop running and can’t respond, the server disk space gets above 80%, or there’s a system security update available. Check these daily and only message me when something needs my attention.”

What this does: Creates background scheduled tasks (cron jobs) that run quietly. You’ll hear nothing most days — which means everything is fine.

6. Set Up Automatic Updates

Keep everything current without thinking about it.

“Check for OpenClaw updates and system security updates every day. When an OpenClaw update is available, tell me what’s new and

ask before installing. For security updates, install them automatically and just let me know what was patched.”

What this does: Your assistant monitors for both OpenClaw feature updates (asks before installing) and OS security patches (installs automatically). You stay current without manual checking.

7. Verify Your Backups Are Working

You set up backup credentials in Part B. Now make sure backups are actually running.

“Check my backup setup. When was the last backup taken? Is the schedule running correctly?”

What this does: Confirms your safety net is in place. If something ever goes wrong, you can restore from a backup.

8. Set a Cost Warning

Avoid surprise bills.

“REMEMBER: Monitor our conversation length. When context reaches 200,000 tokens, warn me so I can decide whether to compact the conversation or keep going. Long conversations cost more because you re-read the entire history with each message.”

Then check your current usage:

```
/status
```

What this does: The REMEMBER sets a permanent alert threshold. The /status command shows your current context usage anytime. Use /compact to shrink a long conversation, or /new to start completely fresh.

9. Teach It the Recovery Procedure

In case you ever need to troubleshoot.

“REMEMBER: If I ever say ‘you were unresponsive’ or ‘Telegram stopped working,’ walk me through the AWS Console Connect recovery steps: sign into AWS, go to EC2, click Instances, click my instance ID, click Connect twice, then type ‘openclaw tui’ to verify you’re running. If the TUI works but Telegram doesn’t, tell me to run ‘openclaw gateway restart’.”

What this does: Your assistant will know how to guide you through recovery even if you forget the steps.

10. Set Up Background Tasks

Your assistant can run tasks on a schedule without you lifting a finger. These are called **cron jobs** — they run in isolated sessions, separate from your conversation, and only message you when something needs your attention.

Steps 5 and 6 above already created your first cron jobs (monitoring and updates). Here are more ideas:

“Check my email every morning at 7 AM and send me a digest of anything important.”

“Remind me every Sunday evening to plan my week.”

You can always ask “What cron jobs do I have running?” to see your full automation schedule. Any recurring task you can describe, your assistant can schedule.

11. Take It for a Spin

Now that everything is configured, try some real tasks to see what your assistant can do.

“Search the web for [something you’re interested in] and give me a summary.”

“What day of the week is my birthday, [date]? Add it to your memory so you remember next year.”

“Write a short thank-you note to [person] for [reason]. Just show me the text.”

“What files are in my workspace? Give me an overview of how things are organized.”

What these do: Gets you comfortable with your assistant’s capabilities — web search, memory, writing, and file management. Try anything you’d ask a human assistant to do.

What’s Next?

You now have a personalized, secured, monitored assistant that knows who you are and how you like to work. Everything you just did — the personality, your profile, the REMEMBER rules, the security, the monitoring — your assistant saved all of it into files that persist across every conversation. These are living documents. You can update them anytime just by talking:

- “I’d like you to be a bit more casual — update your personality”
- “Add to my profile that I’m learning Spanish”
- “REMEMBER my anniversary is June 15th”

From here:

- **Explore Part F** to add optional capabilities (image generation, voice transcription, web search API)
- **Ask your assistant to create skills** for any multi-step process you do repeatedly
- **Just use it** — the best way to learn what it can do is to ask it to do things

If something goes wrong, remember: AWS Console → EC2 → Instances → click your instance → Connect → Connect → type `openclaw tui`.

Part F: Optional API Keys & Services

You already set up the most important API key — Anthropic (Claude) — back in Part B. Everything below is **optional** and adds extra capabilities to your assistant. You can set these up now or come back to them later.

Good news: Just tell your assistant your API keys and it will add them to the configuration for you!

1. Google AI API (For Image Generation)

What it is: Google’s AI can generate images from text descriptions.

Why you want it: Ask your assistant to “generate an image of a sunset over mountains” and it can!

How to get it:

1. Go to aistudio.google.com
2. Sign in with your Google account
3. Accept the terms
4. Click “Get API key” or “Create API key”
5. Copy the key (starts with AIza)

Add it to OpenClaw: Tell your assistant:

“Add my Google AI API key: AIza-xxxxx”

Cost: Free tier is quite generous! Most personal use stays free.

2. OpenAI API (For Voice & Embeddings)

What it is: OpenAI makes ChatGPT, but we’re using them for specific tools:
- Whisper: Transcribes voice messages to text
- Embeddings: Helps with memory search
- GPT (fallback): Backup if Claude is down

Why you want it: Send your assistant voice messages and it’ll understand you!

How to get it:

1. Go to platform.openai.com
2. Sign up or log in
3. Go to “API keys”
4. Click “Create new secret key”
5. Give it a name like “OpenClaw”
6. Copy the key (starts with sk-proj- or sk-)

Add it to OpenClaw: Tell your assistant:

“Add my OpenAI API key: sk-proj-xxxxx”

Cost: Pay-per-use, typically \$2-5/month for voice transcription and embeddings.

3. Brave Search API (For Web Search)

What it is: A privacy-focused search engine API that lets your assistant search the web.

Why you want it: Ask “what’s the weather tomorrow?” or “latest news about [topic]” and your assistant can search for current information.

How to get it:

1. Go to brave.com/search/api
2. Click “Get started”
3. Sign up with your email
4. Verify your email
5. Go to your dashboard
6. Copy your API key

Add it to OpenClaw: Tell your assistant:

“Add my Brave Search API key: xxxxx”

Cost: Free tier includes 2,000 searches per month - plenty for personal use!

Setting Up Dropbox for File Sharing (Optional — Do Later If You Want)

Your assistant can create files (documents, images, spreadsheets, etc.), but those files live on the server — you can’t easily see them on your computer. The solution? Give your assistant access to a Dropbox folder!

You can skip this section for now and set it up later — it’s not required to get started.

Think of it like a shared mailbox: your assistant can PUT files there for you to access, and you can DROP files there for your assistant to work with.

Setting Up Dropbox Access

Already on a Team Dropbox? (like a work or family account) That works great! You don't need a separate account. The access token you create will let your assistant access any folders you have access to. Just pick a folder to use (like "Assistant" or "ClawBot/YourName").

Getting your access token:

1. Go to **dropbox.com/developers** (this is separate from your regular Dropbox — it's where you create apps)
2. Click "Create apps"
3. Choose "Scoped access"
4. Choose "Full Dropbox" access
5. Name your app (e.g., "MyAssistant")
6. Click "Create app"
7. Go to the **Permissions** tab and check these boxes:
 - files.metadata.read
 - files.content.write
 - files.content.read
8. Click "Submit" to save permissions
9. Go back to the **Settings** tab
10. Under "OAuth 2", click "Generate" next to "Access Token"
11. Copy the long token that appears

Tell your assistant:

"I've set up Dropbox for file sharing. My access token is [paste token]. Please set up a Dropbox helper script and save files to the Assistant/ folder."

Your assistant will create the necessary script and configuration automatically!

How to Use It

Once set up, you can say things like: - "Save that report to Dropbox" - "I put a photo in Dropbox — can you look at it?" - "Generate an image and put it in my Dropbox"

Your assistant will save files to the shared folder, and they'll appear on your computer automatically through Dropbox sync!

Tip: Create subfolders to stay organized (e.g., Assistant/Images, Assistant/Documents, Assistant/Projects).

Setting Up a Dedicated Gmail Account (Optional — Do Later If You Want)

One of the most useful things you can do is create a dedicated Gmail account for your assistant. This gives your assistant its own inbox to monitor — perfect for receiving notifications, forwarding emails for review, or having a contact address for services.

You can skip this section for now and set it up later — it's not required to get started.

Why a Dedicated Gmail?

- Forward emails you want your assistant to handle
- Receive notifications from services (AWS alerts, etc.)
- Give your assistant a way to “see” emails without accessing your personal inbox
- Keep assistant-related emails separate and organized

Creating the Gmail Account

1. Go to gmail.com in a private/incognito browser window
2. Click “Create account”
3. Choose a name like “yourname.assistant@gmail.com” or similar
4. Complete the signup process
5. IMPORTANT: Write down the password somewhere safe!

Creating an App Password (for IMAP access)

Your assistant needs a special “app password” to check this inbox — your regular password won't work due to Google's security settings.

1. Go to myaccount.google.com
2. Click “Security” in the left sidebar
3. Under “How you sign in to Google,” click “2-Step Verification”
4. Set up 2-Step Verification if you haven't (required for app passwords)
5. Once enabled, go back to Security
6. Search for “App passwords” or find it under 2-Step Verification
7. Click “App passwords”
8. Select “Mail” and “Other (custom name)”
9. Name it “OpenClaw” and click “Generate”
10. Copy the 16-character password that appears (spaces don't matter)

This app password lets your assistant check the inbox securely.

Setting It Up

Once you have the Gmail account and app password, just tell your assistant:

“I’ve set up a Gmail account for you to monitor. The address is yourname.assistant@gmail.com and the app password is [paste the 16-character password]. Please set up email checking.”

Your assistant will create the necessary script and add the credentials to its configuration.

How to Use It

Once set up: - Forward emails to your assistant’s Gmail for review - Tell your assistant how often to check (e.g., “Check my email every hour” or “Check email twice a day”) - Ask “Check my assistant inbox” anytime for an immediate check - Your assistant will alert you to important messages

Tip: Use Gmail filters to auto-label or forward specific emails. For example, forward all receipts or newsletters to your assistant for organizing.

Part G: Getting the Most Out of Your Assistant

Now that everything’s running, here are some tips for using your assistant effectively.

Sub-Agents: Your Assistant’s Helpers

For big tasks, your assistant can spawn “sub-agents” — separate AI sessions that work in the background. This is like your assistant delegating work to helpers.

Why this matters: - Saves money: Sub-agents have clean, small memory (cheaper!) - Stays responsive: Main conversation isn’t blocked by long tasks - Better results: Each sub-agent focuses on one task

You don’t need to do anything special — your assistant automatically uses sub-agents when appropriate. You might see messages like “I’ve spawned a sub-agent to handle that” followed by a summary when it’s done.

The Thinking Strategy: High Main, Low Sub-Agents

Here’s a cost-saving approach that gives you the best of both worlds:

- **Main session:** Set `thinkingDefault: high` in your config. This gives you careful, thorough reasoning for your direct conversations and important decisions.
- **Sub-agents:** Your assistant dispatches these with lower or no thinking, since they’re handling routine, well-defined tasks.

Why this works: Your main session handles strategy, judgment calls, and complex questions (where deep thinking matters). Sub-agents handle mechanical tasks like “update this WordPress page” or “compile this report” (where speed matters more than deliberation).

To set this up, just tell your assistant:

“Set your default thinking level to high”

Your assistant will update the configuration. It already knows to spawn sub-agents with lower thinking levels for routine work, so this keeps your costs reasonable while maintaining high-quality reasoning where it counts.

Understanding Your Assistant's Automation Architecture

Your assistant runs multiple types of sessions behind the scenes, each with a specific purpose. Understanding this helps you make the most of automation without confusion:

Session Types:

Main session (your conversation): - This is where you chat with your assistant - Full memory and context - Interactive and conversational - Can spawn sub-agents for complex tasks

Cron sessions (scheduled automation): - Isolated, independent sessions - Run on exact schedules - Alert you only when needed - Examples: backups, monitoring, digests, social media posts

Sub-agent sessions (delegated work): - Spawned by main session for complex tasks - Clean, task-focused memory - Report back when complete - Examples: “research this topic,” “compile that report”

Real-World Example:

Here's how Steve's assistant architecture works—this shows you what's possible:

Cron jobs (isolated sessions): - **Daily at 6:00 AM:** Email digest compilation (checks inbox, categorizes messages, sends morning briefing) - **Daily at 11:55 PM:** System health monitoring (checks server status, disk space, memory usage) - **Every 3 days at 9:00 AM:** BrainStream Facebook post (selects content, generates post, publishes) - **Sundays at 1:00 AM:** Automatic server backups (creates AWS snapshot, deletes old backups) - **Mondays at 8:00 AM:** Weekly summary (compiles what happened last week)

Heartbeats (shared session - Steve uses very few): - Checking calendar for conflicts during active work hours - Following up on ongoing projects we discussed recently

Sub-agents (spawned as needed): - “Research OpenClaw features and write a comparison guide” - “Update all three versions of the installation documentation” - “Compile analytics from last month's newsletter campaigns”

Notice the pattern: **Scheduled/recurring = cron jobs. One-time complex tasks = sub-agents. Opportunistic checks = rare heartbeats.**

Why This Architecture Matters:

Cost efficiency: Cron jobs have clean memory context (cheaper), main session has full context (more expensive). Using the right tool for each job keeps costs reasonable.

No interruptions: Cron jobs run silently and only alert you when something needs attention. Your main conversation stays focused.

Reliability: Cron jobs run on exact schedules, regardless of whether you're actively chatting. Your backups, monitoring, and automation happen like clockwork.

Flexibility: You can add, modify, or remove cron jobs anytime. "Add a cron job to check my website uptime every hour" or "remove the Facebook posting cron job" — just ask.

Setting Up Your Own Automation:

Start simple and build over time:

Week 1: Just use your assistant conversationally—no automation needed yet.

Week 2: Add one cron job for something you check manually every day:

"Set up a daily cron job to check my email at 7 AM and send me a digest"

Week 3: Add automation for a recurring task:

"Create a weekly backup cron job every Sunday at 1 AM"

Month 2: Expand based on your patterns:

"I notice I check my server status every morning—can you automate that as a cron job?"

Your assistant will guide you through setting up each automation and help you decide whether something should be a cron job, heartbeat, or neither.

The key insight: You don't need to understand all this upfront. Start with conversations, add automation when you see repetitive tasks, and let your assistant recommend the right approach.

Managing Your Context Window (Important for Costs!)

Your assistant has a **context window** — think of this as its short-term memory for the current conversation. Everything you've discussed in the

current session lives in this buffer. The default size (200,000 tokens) works well for most people and keeps costs reasonable.

After about 200,000 tokens, the per-message cost increases significantly. Your assistant will warn you when the context is getting high. You can increase the context window up to 1 million tokens if you need longer conversations, but for most people the default is the sweet spot.

When you get a warning that context is high, you have two choices:

- **Continue as-is** — if you're in the middle of something complex and don't want to lose any context, just keep going. The higher cost is worth it when continuity matters.
- **Use the /compact command** — this summarizes your conversation history, freeing up the context window while preserving the key points. Think of it like your assistant taking notes and then clearing its desk.

The best time to /compact is **when you're naturally shifting topics** — if you just finished a big project discussion and want to move on to something completely different, that's the perfect moment. Your assistant will compress everything into a summary, and you'll be back to a clean, low-cost context.

Image Generation Capabilities

If you set up the Google AI API (Nano Banana Pro), your assistant can do more than just generate images:

- Generate images from descriptions (“a cozy coffee shop at sunset”)
- Edit existing images (send a photo and ask for changes)
- Remove objects from photos (“remove the person in the background”)
- Add objects to photos (“add a cat sitting on the chair”)
- Change backgrounds (“put me on a beach instead”)
- Multiple resolutions (1K, 2K, or 4K quality)

Just describe what you want naturally — your assistant figures out which capability to use!

Adding Other Telegram Users

Want to give family members or colleagues access to your assistant? Here's how:

1. Have them install Telegram
2. Have them message **@userinfobot** on Telegram
 - This bot will reply with their user ID (a number like 123456789)

- Have them send you this number
3. Tell your assistant:

“Add Telegram user [their ID number] to my allowed users”

Your assistant will update the configuration and restart itself.

Now they can message your bot too! Note: They’ll share the same assistant and memory, so only add people you trust.

Part H: What Else Can Your Assistant Do?

Congratulations! You have a working personal AI assistant! 🎉

Here are some things to explore as you get comfortable:

Setting Up Automatic Backups

Remember those AWS backup credentials you saved in Part B? Now it’s time to use them.

Tell your assistant:

“Set up automatic weekly backups of my server. My AWS Access Key ID is [paste first key] and my Secret Access Key is [paste second key]. Run backups Sunday at 1 AM Eastern time, and keep only the two most recent backups.”

Your assistant will set up a weekly backup schedule. Each week, it creates a complete snapshot of your server, and automatically deletes old backups so you don’t accumulate charges.

Keeping Things Updated

Your OpenClaw installation and server need regular updates to stay secure and get new features. Here’s what you need to know:

Why Updates Matter

Think of updates in two categories:

Security patches: These fix vulnerabilities that hackers could exploit. Ignoring them is like leaving your front door unlocked.

New features and bug fixes: OpenClaw gets improvements regularly — better performance, new capabilities, fixes for things that weren’t working quite right.

Keeping everything updated means your assistant stays secure, runs smoothly, and gets access to the latest features.

Two Types of Updates

1. OpenClaw Updates (New Features & Improvements)

OpenClaw itself gets updated regularly. To check if a new version is available:

- See what version is installed: `openclaw --version`
- Check the latest available: `npm show openclaw version`

If those numbers are different, an update is available!

Your assistant can handle the entire update process for you. It will: - Check for updates - Show you what's new in the release - Install the update when you approve it - Restart itself automatically

2. System Security Updates (Operating System Patches)

Your Ubuntu server also needs security patches. These are separate from OpenClaw and protect the underlying computer.

To check for security updates:

```
sudo apt update && apt list --upgradable | grep security
```

This shows any security patches waiting to be installed.

Again, your assistant can handle this — it will tell you what needs updating and install it when you approve.

The Easy Way: Let Your Assistant Handle It

Instead of manually checking, just tell your assistant:

“Check daily for OpenClaw updates and system security updates. Let me know when updates are available, and ask me before installing anything.”

Your assistant will add this to its daily routine. Each day, it checks both OpenClaw and system security updates. When something is available, it will message you with details and wait for your approval before installing.

The update process looks like this:

1. Your assistant: “A new OpenClaw version is available (v2026.2.1). Changes: improved memory handling, faster search, bug fixes. Should I install it?”
2. You: “Yes, go ahead”
3. Your assistant: “Installing now... Done! I’m restarting to apply the update. Back in a moment!”

That's it! No Terminal commands, no technical knowledge needed — just approve and your assistant does the work.

Important: Your assistant will always ask before installing anything. Updates occasionally require a restart (especially OpenClaw updates), so you'll have a brief moment where your assistant is offline. Your assistant will warn you before this happens.

Keeping Your Server Secure

Your AWS server is already more secure than most setups out of the box. Let's make sure you understand why, and add a few easy improvements.

What AWS Already Does for You

When you set up your server in Part B, AWS automatically gave you several security features:

SSH Key Authentication: Remember that .pem key file? That's your only way to log into the server. There's no password to guess — if someone doesn't have your key file, they can't get in. This is much more secure than a username and password.

Security Groups (Your Firewall): AWS has a built-in firewall called a Security Group that controls which network traffic can reach your server. By default, only SSH (port 22) is open — everything else is blocked. Your assistant communicates through Telegram's cloud servers, so it doesn't need any additional ports open to the internet.

Non-Root User: You're logging in as ubuntu, not as root (the all-powerful administrator account). This means even if someone somehow got access, they couldn't immediately do the most dangerous things without an extra step.

Isolated Server: OpenClaw runs on its own dedicated server in the cloud, completely separate from your personal computer. If anything ever went wrong, your personal files, photos, and data are untouched.

Easy Security Improvements

Here are three things you can do in about five minutes that significantly improve your security:

1. Install Fail2ban (Stops Brute Force Attacks)

Fail2ban watches for repeated failed login attempts and automatically bans the attacker's IP address. Think of it as a bouncer who kicks out anyone who keeps trying wrong passwords.

Tell your assistant:

“Install fail2ban and configure it to ban IPs after 3 failed SSH attempts for 24 hours.”

Your assistant will install and configure it. Once running, it works silently in the background — you’ll never need to think about it again.

2. Ask Your Assistant to Audit Its Own Security

This is one of the coolest things about having an AI assistant — it can check its own security! Just ask:

“Audit your own security setup and tell me if anything needs attention.”

Your assistant will check things like: whether SSH is properly configured, whether unnecessary ports are open, whether your system is up to date, and whether there are any obvious vulnerabilities. It’s like having a security consultant on call 24/7.

3. Set Up Real-Time Alerts

Your assistant can monitor itself and message you when something is wrong:

“Set up monitoring to alert me if: the server is running low on disk space, the server hasn’t been backed up recently, OpenClaw stops running, or anything unusual happens.”

Your assistant will create background checks (cron jobs) that run periodically and only message you when something actually needs your attention. Most days, you’ll hear nothing — which means everything is fine.

For Advanced Users

If you want to go further, here are some additional hardening options. These aren’t necessary for most people, but they’re available:

- **Tailscale VPN:** Creates a private network between your devices. Useful if you want to access your server’s web interface directly without exposing any ports to the internet. Free for personal use. See tailscale.com.
- **UFW Firewall:** An additional firewall layer on the server itself, on top of AWS Security Groups. Belt-and-suspenders approach.
- **Docker Sandboxing:** Runs sub-agent tasks in isolated containers so they can’t access sensitive files. Significant setup effort, but worth it for high-security environments.

For most users, the combination of AWS Security Groups + SSH keys + Fail2ban + your assistant’s self-monitoring covers everything you need.

Other Messaging Channels

OpenClaw can also connect to WhatsApp, Discord, and SMS. Check the OpenClaw documentation when you're ready to explore these.

Browser Automation (Advanced — Optional)

OpenClaw has the ability to control a web browser on the server using a tool called Playwright. This lets your assistant log into websites, fill out forms, and take screenshots of web pages.

Most people won't need this. Your assistant can already search the web, fetch web pages, and interact with services that have APIs — all without a browser. Browser automation is mainly useful for sites that require a visual login session or don't have any other way to access their data.

If you're interested, just tell your assistant:

“Install Playwright browser automation with Chromium.”

It takes a few minutes and about 500MB of disk space. But honestly, start without it and add it later only if you find a specific need.

Skills, Habits, and Knowing Which to Use

As you use your assistant more, you'll develop processes you want done the same way every time. There are two good ways to lock these down, and knowing which to use matters:

Skills are best for complex, multi-step processes where consistency is critical. Think of them as detailed recipe cards your assistant follows precisely. For example, compiling a monthly newsletter involves selecting content, generating a writing tip, rendering an email template, sending you a preview, and archiving used items. That's a lot of steps with specific rules — perfect for a skill. Your assistant can create skills for you:

“Let's create a skill for our newsletter process. It should document every step so it's done the same way each month.”

Behavioral prompts in AGENTS.md are better for ongoing habits that should happen throughout the day without being told. For example, maintaining a daily status report works best as a standing instruction: “After completing any significant task, append it to today's status report immediately.” Your assistant reads AGENTS.md every session, so these become automatic habits rather than one-time procedures.

A good rule of thumb:

- If it's a process with a clear start and end (compile a newsletter, process a batch of data, update a document) → **make it a skill**

- If it's an ongoing behavior or daily habit (keep a status report, check email, monitor something) → **add it to AGENTS.md** with a nightly organization task

Both approaches prevent “drift” — where your assistant gradually starts doing things differently over time. The key insight is: **write it down, don't rely on memory**. Your assistant's memory can be compacted or lost between sessions, but skills and AGENTS.md persist forever.

Browse the OpenClaw documentation or ask your assistant to help you identify which of your processes would benefit from being captured as skills.

Conclusion

You did it! You now have your own personal AI assistant running in the cloud, customized to your personality and needs.

Remember: - Your assistant learns about you over time; the more you use it, the better it gets. - Use the REMEMBER convention to teach it important facts. - Need help? Ask your assistant for help (seriously, it can troubleshoot itself!)

Welcome to the world of personal AI assistants. Enjoy! 🚀

Really Technical Stuff (You Probably Don't Need This)

This section covers common issues and how to fix them. Most people will never need this, but it's here just in case.

Issue 1: Assistant Not Responding (Most Common!)

Problem: You send messages in Telegram, but your assistant doesn't reply.

Possible causes: - OpenClaw stopped running or needs a restart. - Your server is offline. - You ran out of API credits.

How to fix — the AWS Console method (no Terminal required!):

This is the easiest way to check on your server and fix most problems. You don't need Terminal, SSH, or any special software — just a web browser.

1. Go to **aws.amazon.com** and sign in to your AWS account.
2. In the search bar at the top, type **EC2** and click on it.
3. Click **Instances** in the left sidebar.
4. You should see your server listed. First, check: is the **Instance state** showing “Running” with a green dot? If not, select it and click **Start instance**, then wait a minute or two.
5. If the instance IS running, click on the **Instance ID** (the blue link that looks like i-0abc123def456).

6. Click the **Connect** button near the top of the page.
7. On the next screen, leave the defaults and click **Connect** again.
8. A black terminal window will open right in your browser! You're now connected to your server.
9. Type: `openclaw tui`
10. If the TUI opens and your assistant responds, try messaging it in Telegram again — it may just need a moment to reconnect.
11. If the TUI shows an error, or if your assistant still isn't responding in Telegram, type: `openclaw gateway restart` Wait about 30 seconds, then try Telegram again.
12. If none of that works, try a full restart: `sudo systemctl restart openclaw-gateway`

Still stuck? Check your API credits at console.anthropic.com. If your balance is zero, your assistant can't think — add more credits and it will come back to life.

Alternative: SSH from Terminal (if you prefer the traditional method)

If you're comfortable with Terminal/PowerShell, you can also SSH in directly:

```
ssh -i ~/Downloads/your-key-name.pem ubuntu@your-server-ip
```

Then run the same commands: `openclaw tui` or `openclaw gateway restart`.

Issue 2: “Cooldown” Warnings or Rate Limiting

Problem: Your assistant says something about “cooldown,” “rate limit,” or temporarily can't respond. It might suggest adding a different AI provider as a fallback.

What's happening: Anthropic limits how fast you can send messages based on your credit balance. New accounts with small balances (\$5-10) have very low limits. During setup, when you're chatting frequently, it's easy to hit them.

How to fix:

1. Go to **console.anthropic.com** → **Billing**
2. Add credits to bring your balance to at least **\$50**
3. Set auto-reload to trigger at **\$20** so you never run low

The rate limits increase automatically as your balance grows. With \$50 in credits, you'll rarely see cooldown warnings.

⚠ Important: If your assistant suggests switching to a different AI provider (like OpenAI) as a fallback, **don't do it** during setup. The assistant behaves differently on other models and may not be able to switch itself

back properly. Instead, just wait a minute for the cooldown to pass, or add more credits. If your assistant has already gotten into a confused state from switching providers, the safest fix is to re-run the setup wizard:

```
openclaw setup
```

This regenerates a clean configuration. Walk through the wizard again with your Anthropic key and you'll be back to normal.

Issue 3: Messages Delayed or Slow

Problem: Your assistant responds, but it takes a long time.

Possible causes: - Complex request (web search, browser automation). - Server overloaded. - Network latency.

How to fix: This is usually normal for complex tasks. If *everything* is slow, check your server's resources (CPU/Memory) via SSH. If usage is constantly at 100%, consider upgrading your instance size.

Issue 4: API Key Errors

Problem: Your assistant says "API key invalid" or "API authentication failed."

How to fix: Your key might have expired. Get a new one from console.anthropic.com, then tell your assistant: > "Update my Anthropic API key: sk-ant-..."

Issue 5: Out of Disk Space

Problem: Your assistant says "No space left on device" or operations fail mysteriously.

How to fix: SSH into your server and check disk usage (`df -h`). If usage is above 90%: 1. **Clean up:** Ask your assistant to "Clean up old logs and temporary files." 2. **Expand storage:** In AWS Console (EC2 → Volumes → Modify Volume), increase the size, then tell your assistant to expand the filesystem.

Issue 6: Server Can't Be Reached via SSH

Problem: Connection refused or timed out.

Possible causes: - Server stopped. - IP address changed (if you rebooted). - Security group rules (port 22). - Key permissions.

How to fix: 1. **Check AWS:** Is the instance running? 2. **Check IP:** Did the Public IPv4 address change? Use the new one. 3. **Check Security Group:** Ensure port 22 is open to your IP or "Anywhere". 4. **Check Key Permissions:** Make sure `chmod 400 your-key.pem` was run (Mac/Linux).

Best Practices to Avoid Issues

- Keep OpenClaw and your system updated.
- Monitor costs and disk space regularly.
- Set up automatic backups.
- Don't install random software from untrusted sources.

Webhooks

External services can trigger your assistant automatically. For example: - Teachable can alert you when students enroll. - Stripe can notify you when payments come through. - Your website contact form can forward inquiries.

See the OpenClaw documentation at docs.openclaw.ai for setup details. Just ask your assistant to help you set these up. ### Real-Time Email Monitoring

By default, your assistant can check email via scheduled cron jobs (e.g., every hour, or twice daily). That works fine for most people, but what if you want your assistant to respond to your emails instantly?

Enter **IMAP IDLE** — a way for the email server to tap your assistant on the shoulder the instant a new email arrives, instead of your assistant having to keep checking. Think of it like this: instead of walking to your mailbox every half hour to see if you have mail, the mail carrier rings your doorbell the moment a letter arrives.

Why You'd Want This

- **Instant responses:** Your assistant sees your email within seconds and can reply or take action immediately
- **No waiting:** Instead of "I'll check on the next scheduled cron run," your assistant is already on it
- **More natural:** Email conversations feel more like real conversations

How It Works

A Python script maintains a persistent connection to Gmail using IMAP IDLE. When a new email arrives from your address, the email server notifies the script immediately, and the script triggers an OpenClaw wake event — basically telling your assistant "Hey, you have mail from the owner, go check it now!"

The script runs as a systemd service, which means it starts automatically when your server boots and restarts itself if it ever crashes.

Prerequisites

You need a Gmail account for your assistant with an app password set up. We covered this in **Part F: Setting Up a Dedicated Gmail Account** — if you haven't done that yet, go back and set it up first.

The Script

Here's the complete, working email watcher script. Your assistant can help you set this up, but if you're doing it yourself, save this as `email-watcher.py` somewhere in your workspace (like `~/clawd/scripts/`):

```
#!/usr/bin/env python3
"""
IMAP IDLE Email Watcher
Monitors a Gmail inbox and triggers OpenClaw wake events for emails
from the owner.
"""

import imaplib
import email
import select
import socket
import time
import sys
import re
import subprocess
from datetime import datetime

# ===== CONFIGURATION =====
# Your assistant's Gmail address and app password
IMAP_HOST = "imap.gmail.com"
IMAP_USER = "your-assistant@gmail.com"
IMAP_PASSWORD = "xxxx xxxx xxxx xxxx" # Gmail app password

# Your email addresses (the assistant watches for emails from you)
OWNER_ADDRESSES = [
    "you@example.com",
    "you@work.com"
]

# Path to openclaw (run 'which openclaw' to find yours)
OPENCLAW_PATH = "/usr/local/bin/openclaw"

IDLE_TIMEOUT = 1740 # 29 minutes (Gmail drops IDLE connections after
~29 min)
MAX_BACKOFF = 300 # Max reconnection delay (5 minutes)
# =====

def log(message):
```

```

timestamp = datetime.now().strftime("%Y-%m-%d %H:%M:%S")
print(f"[{timestamp}] {message}", flush=True)

def extract_email_address(from_header):
    if not from_header:
        return None
    match = re.search(r'<([^\s@]+)>|([^\s<>]+@[^\s<>]+)', from_header)
    if match:
        return (match.group(1) or match.group(2)).lower()
    return None

def is_from_owner(from_address):
    if not from_address:
        return False
    return any(addr.lower() in from_address.lower() for addr in
OWNER_ADDRESSES)

def trigger_wake_event(subject):
    try:
        message = f"New email from owner: {subject}. Check inbox and
respond."
        cmd = [OPENCLAW_PATH, "system", "event", "--mode", "now", "--
text", message]
        log(f"Triggering wake event: {message}")
        result = subprocess.run(cmd, capture_output=True, text=True,
timeout=10)
        if result.returncode == 0:
            log("Wake event triggered successfully")
        else:
            log(f"Wake event failed: {result.stderr}")
    except Exception as e:
        log(f"Error triggering wake event: {e}")

def fetch_email_details(mail, email_id):
    try:
        status, data = mail.fetch(email_id, '(RFC822.HEADER)')
        if status != 'OK':
            return None, None
        msg = email.message_from_bytes(data[0][1])
        return msg.get('From', ''), msg.get('Subject', '(no subject)')
    except Exception as e:
        log(f"Error fetching email details: {e}")
        return None, None

def connect_imap():
    log(f"Connecting to {IMAP_HOST}...")
    mail = imaplib.IMAP4_SSL(IMAP_HOST)
    mail.login(IMAP_USER, IMAP_PASSWORD)
    log("Connected and authenticated successfully")
    return mail

```

```

def idle_wait(mail, timeout):
    tag = mail._new_tag().decode()
    mail.send(f'{tag} IDLE\r\n'.encode())

    sock = mail.socket()
    old_timeout = sock.gettimeout()
    sock.settimeout(30)
    try:
        response = mail.readline()
        if b'+' not in response:
            log(f"IDLE not accepted:
{response.decode(errors='replace').strip()}")
            return 'error'
    except Exception as e:
        log(f"Error reading IDLE response: {e}")
        return 'error'
    finally:
        sock.settimeout(None)

    log("IDLE mode active, waiting for new emails...")

    result = 'timeout'
    try:
        readable, _, _ = select.select([sock], [], [], timeout)
        if readable:
            sock.settimeout(5)
            try:
                data = sock.recv(4096)
                decoded = data.decode(errors='replace')
                log(f"IDLE notification: {decoded.strip()}")
                if 'EXISTS' in decoded:
                    result = 'new_mail'
            except Exception as e:
                log(f"Error reading IDLE data: {e}")
                result = 'error'
        else:
            log("IDLE timeout reached (29 min), re-entering IDLE")
    except Exception as e:
        log(f"Error in select(): {e}")
        result = 'error'

    try:
        sock.settimeout(30)
        mail.send(b'DONE\r\n')
        for _ in range(10):
            resp = mail.readline()
            if resp and tag.encode() in resp:
                break
    except Exception as e:

```

```

        log(f"Error exiting IDLE: {e}")
        result = 'error'
    finally:
        sock.settimeout(old_timeout)

    return result

def process_new_emails(mail):
    try:
        mail.select("INBOX")
        status, messages = mail.search(None, 'UNSEEN')
        if status != 'OK' or not messages[0]:
            return
        for email_id in messages[0].split():
            from_header, subject = fetch_email_details(mail, email_id)
            if from_header:
                from_address = extract_email_address(from_header)
                if is_from_owner(from_address):
                    log(f"Email from owner: {subject}")
                    trigger_wake_event(subject)
    except Exception as e:
        log(f"Error processing emails: {e}")

def main():
    log("Email watcher starting...")
    backoff = 1
    while True:
        mail = None
        try:
            mail = connect_imap()
            mail.select("INBOX")
            backoff = 1
            while True:
                result = idle_wait(mail, IDLE_TIMEOUT)
                if result == 'new_mail':
                    process_new_emails(mail)
                elif result == 'error':
                    break
        except Exception as e:
            log(f"Connection error: {e}")
        if mail:
            try: mail.close(); mail.logout()
            except: pass
        log(f"Reconnecting in {backoff}s...")
        time.sleep(backoff)
        backoff = min(backoff * 2, MAX_BACKOFF)

if __name__ == "__main__":
    try:
        main()

```

```
except KeyboardInterrupt:
    log("Stopped")
    sys.exit(0)
```

Important technical note: This script uses `select.select()` for waiting instead of socket timeouts. This matters! Using `socket.setdefaulttimeout()` will cause the IDLE connection to drop after the timeout period instead of waiting the full 29 minutes. If your assistant writes this script from scratch, make sure it uses `select.select()`.

The Systemd Service

To run the script automatically as a background service, create a systemd service file. Save this as `/etc/systemd/system/email-watcher.service`:

```
[Unit]
Description=Email Watcher for OpenClaw
After=network.target

[Service]
Type=simple
User=ubuntu
ExecStart=/home/ubuntu/your-workspace/venv/bin/python3
/home/ubuntu/your-workspace/scripts/email-watcher.py
Restart=always
RestartSec=10

[Install]
WantedBy=multi-user.target
```

Adjust the paths to match your setup — use which `python3` to find your Python path if you're not using a virtual environment.

Then enable and start the service:

```
sudo systemctl daemon-reload
sudo systemctl enable email-watcher
sudo systemctl start email-watcher
```

Check that it's running:

```
sudo systemctl status email-watcher
```

The Easy Way: Ask Your Assistant

Instead of manually setting all this up, just give your assistant this prompt:

"I'd like you to set up real-time email monitoring. Here's a Python script that watches my Gmail inbox using IMAP IDLE and triggers a wake event when I send you an email. Please save this as a script,

set it up as a systemd service so it runs automatically, and test it by having me send you a test email.”

Then paste the script above (with your real credentials filled in), and your assistant will handle the setup, testing, and troubleshooting.

Testing: Once it's running, send a test email to your assistant's Gmail address. You should see your assistant respond within a few seconds instead of waiting for the next scheduled check!

Last updated: February 2026 **OpenClaw version:** 2026.x