

Table Of Contents

[Before You Start](#)

[1. API Basics](#)

[2. A Backend Framework](#)

[3. Databases](#)

[4. Authentication and Authorization](#)

[5. Input Validation and Error Handling](#)

[6. File Storage and External Services](#)

[7. Background Jobs and Queues](#)

[8. Caching](#)

[9. Testing](#)

[10. Deployment and Observability](#)

[What's Next?](#)

The Ultimate Backend Roadmap

This is the path I would tell a friend to follow if they wanted to actually become a backend engineer and get a job doing it.

The whole roadmap is built around one idea: you learn backend by shipping backends. Every section ends with a small addition to a single project, a habit tracker API (basic I know, but it's a good way to get started). By the end you will have one complete, deployed, tested backend that touches almost everything a real production service does. It's not going to get you a job, but you'll learn the fundamentals you need to build a larger scale project that will.

Before You Start

This roadmap assumes you can already code. Not well, just the basics. If you do not know what a for loop is, what a function is, or what an if statement does, stop here and go learn one language first.

Pick Python or Java and learn that before anything else.

- **Python** is the fastest way in. Clean syntax, huge ecosystem, and the framework we use in this guide (FastAPI) is Python. It's also the best language to interview in most of the time.
- **Java** is harder up front but it is everywhere in big tech, will teach you really good OOP fundamentals, and Spring Boot is one of the most in-demand backend skills on the market.

1. The Basics

Before you touch a framework or build a system, you need to understand the core building blocks of what you are actually building.

What is a backend?

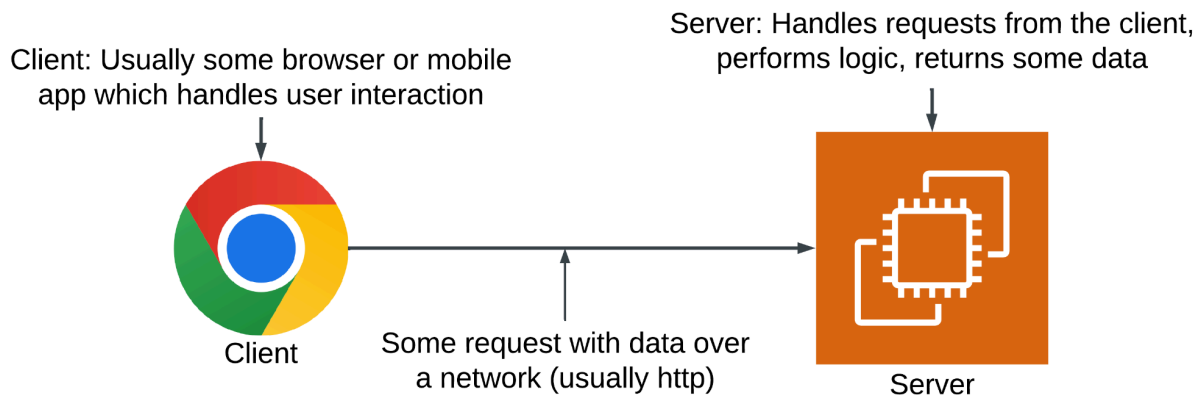
When you use an app, there are two halves. The **frontend** is what you see and tap: the buttons, the screens, the colors. The **backend** is everything you don't see: where the data lives, the logic that decides what happens, and the rules that keep it secure. When you like a post, the frontend shows the heart turn red, but the backend is what actually records that you liked it, makes sure you can't like it twice, and remembers it forever.

As a backend engineer, you build that second half. You do not worry much about how things look. You worry about data, logic, and making sure the whole thing keeps working when a million people use it at once.

What is a server?

A **server** is just a computer that is always on, waiting for other computers to ask it for things. That's it. The laptop in front of you could be a server. A giant machine in an Amazon data center is a server. The only thing that makes a computer a "server" is that its job is to sit there and respond to requests.

The computer doing the asking is called the **client**. Your phone is a client. Your browser is a client. When you open Instagram, your phone (the client) asks Instagram's servers for your feed, and the servers send it back.



This back-and-forth is the entire foundation of the web. 99% of apps you have ever used are, underneath, a client asking a server for something and the server answering. Your whole job as a backend engineer is to build the thing that answers.

What is an API?

An **API** (Application Programming Interface) is the menu of things a server is willing to do. Think of a restaurant. You don't walk into the kitchen and cook your own food. You look at a menu, you order, and the kitchen makes it for you. You don't need to know how the kitchen works, you just need to know what you're allowed to order and how to ask.

An API is that menu. It is the list of requests a server will accept and what it will give you back. When your phone wants your Instagram feed, it doesn't reach into Instagram's database directly

(that would be a disaster). It asks the API: "give me this user's feed," and the API hands it back in a clean, predictable format.

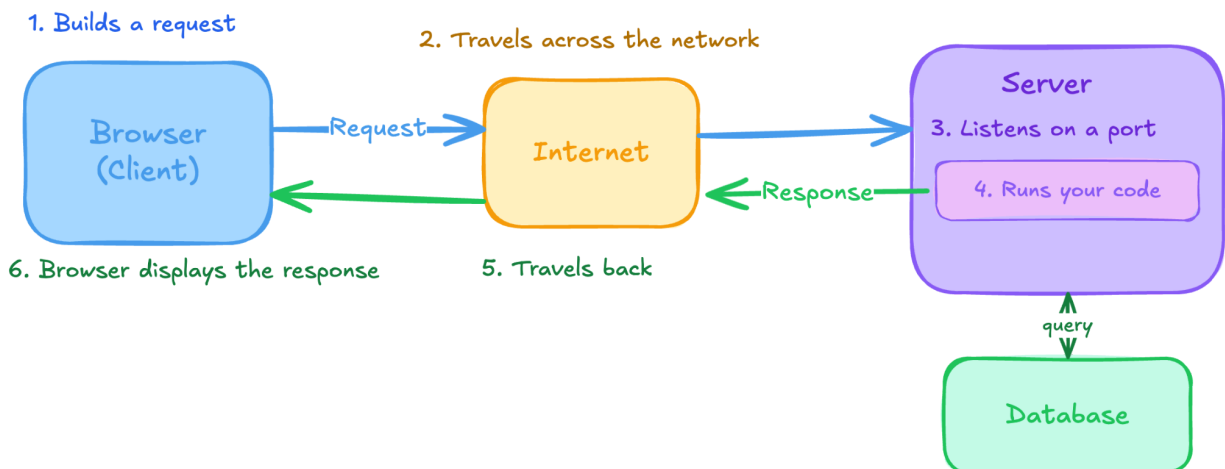
So when we say you're going to "build an API," we mean you're going to build that menu: the set of things your server knows how to do, and the rules for how to ask for them.

How a request actually travels

Let's follow one request from start to finish:

1. The client (say, your browser) wants something. It builds a **request**: a URL saying where to go, a method saying what it wants to do, and maybe some data.
2. That request travels across the internet to the server's address.
3. The server is **listening** on a specific door, called a **port**. It receives the request.
4. The server runs your code to figure out what to do (look something up in the database, check a permission, whatever).
5. The server builds a **response** and sends it back across the internet to the client.
6. The client gets the response and does something with it (shows it on screen).

The Life of a Request



You will hear two words constantly while you build this on your own machine:

- **localhost**: a nickname your computer uses to mean "this same computer." When you run a server on your laptop and visit `localhost`, you are the client and the server, both on one machine. This is how you'll test everything before it ever goes online.

- **Port:** a numbered door on a computer. One machine can run many servers at once, so each one listens on its own port. You'll see things like `localhost:8000` constantly. That just means "the server running on this computer, behind door number 8000."

That's the whole physical picture. A client builds a request, it travels to a server listening on a port, the server runs your code and sends back a response. Now let's look at what's actually inside a request.

HTTP: the language requests are (usually) written in

HTTP is the set of rules that clients and servers use to talk to each other on the web. Every request has a few parts, and you need to understand each one.

HTTP methods

The **method** says what you are trying to do. There are a handful, and they map cleanly onto the four things you ever do with data: create it, read it, update it, delete it (developers call this **CRUD**).

Method	What it does	CRUD	Example
GET	Read data	Read	Get a user's habits
POST	Create new data	Create	Add a new habit
PUT	Replace data entirely	Update	Replace a whole habit record
PATCH	Update part of data	Update	Mark today's habit complete
DELETE	Remove data	Delete	Delete a habit

You will use GET and POST constantly, the others less often. The important idea: the method tells the server your intent. A GET should never change anything, it just reads. A DELETE removes. Keeping these honest is the start of good API design.

URLs and paths

The **URL** is the address of what you want. Break one down:

<https://myapp.com/habits/42>

- <https://> the protocol (HTTP, secured)

- `myapp.com` the server's address
- `/habits/42` the **path**, which points at a specific resource (habit number 42)

The path is where your API's structure lives. `/habits` means "all habits." `/habits/42` means "the one specific habit with ID 42." You design these.

Status codes

When the server responds, it includes a **status code**: a three-digit number that says how it went. You do not memorize all of them. You learn the ranges and a handful of common ones.

Range	Meaning	The ones you'll actually use
2xx	Success	200 OK, 201 Created
3xx	Go look somewhere else	301, 304
4xx	The client messed up	400 Bad Request, 401 Unauthorized, 403 Forbidden, 404 Not Found
5xx	The server messed up	500 Internal Server Error, 503 Unavailable

TLDR: **4xx is your fault as the caller, 5xx is the server's fault.** A 404 means you asked for something that isn't there. A 500 means the server crashed trying to handle your request.

Headers

Headers are extra information attached to every request and response, separate from the main data. They carry context. The big ones you'll meet:

- **Content-Type:** `application/json` tells the other side what format the data is in
- **Authorization:** `Bearer <token>` proves who you are (you'll build this in the auth section)

You don't memorize headers. You just need to know they exist and that they carry the metadata around the actual content.

The body and JSON

The **body** is the actual data being sent. When you create a habit, the details of that habit ride along in the request body. When the server sends your habits back, they ride in the response body.

That data is almost always formatted as **JSON** (JavaScript Object Notation). JSON is just keys and values, and it's readable by humans and machines alike:

```
{  
  
  "habit": "drink water",  
  "completed_today": true,  
  "streak": 12  
}
```

Curly braces hold an object, each key (like "**streak**") points to a value (like **12**). That's the whole format. Almost every modern API speaks JSON, so you'll be reading and writing this constantly.

Putting it together: REST API design

Now you know the pieces. **REST** is just a popular, sensible style for arranging them into a clean API. You do not need theory, you need two rules:

1. **Organize your API around resources (nouns), not actions (verbs).** Your habit tracker has habits, so the resource is `/habits`. You do NOT make URLs like `/getHabits` or `/createNewHabit`.
2. **Let the HTTP method say what you're doing.** The URL names the resource, the method names the action.

So a well-designed habit tracker API looks like this:

Method + Path	What it does
GET <code>/habits</code>	List all habits
POST <code>/habits</code>	Create a new habit
GET <code>/habits/42</code>	Get habit 42
PATCH <code>/habits/42</code>	Update habit 42

Method + Path	What it does
DELETE <code>/habits/42</code>	Delete habit 42

Notice the same path, `/habits/42`, does completely different things depending on the method. That is REST. Clean, predictable, and every backend engineer instantly understands it. This is all the API design you need to start.

Project: Build a habit tracker API

Build a [simple API with FastAPI](#) for tracking habits. You don't need a database yet (we'll add that later) just keep the habits in a Python list in memory. Add endpoints to create a habit (POST `/habits`), list all habits (GET `/habits`), and mark one complete (PATCH `/habits/{id}`).

The real goal here is not the code, it is the wiring. Run the server on your machine, then hit your own API from your computer. FastAPI gives you an interactive `/docs` page automatically, so go to `localhost:8000/docs` in your browser and click buttons to send real requests. Or use `curl` or Postman if you want to feel like a hacker.

Once you can ping your own API and get JSON back, you've done more than WAY too many recent grads I've met, and you can start to build larger systems.

2. A Backend Framework

99% of companies run on a **backend framework**: a toolkit that handles all the repetitive parts (routing requests, parsing JSON, sending responses) so you can focus on your actual logic.

These are the core frameworks I think are worth considering to start with:

Stack	Language	Where you see it
Spring Boot	Java	Big tech, large enterprises
FastAPI	Python	Startups, AI companies, data teams
Express	Node.js	Startups, full-stack JS shops


Spring is one of the best frameworks you can learn if you are aiming for big tech. A lot of startups run Python, and most of the rest run Node. Pick one and go deep. Do not learn all three.

Here is the single most important thing in this whole section: **do not watch twelve tutorials and ship zero backends**. Pick one framework and learn it really well by building something with it. The concepts you learn here transfer to every other framework on earth.

Those transferable concepts are:

- **Routing**: mapping a URL and method (like GET `/habits`) to a specific function in your code
- **Middleware**: code that runs before or after your handler, on every request (logging, auth checks, etc.)
- **Validation**: making sure an incoming request actually has the right shape before you trust it
- **Error handling**: catching failures and returning a clean response instead of letting the server crash
- **Environment variables**: keeping config and secrets (passwords, API keys) out of your code

The Request Lifecycle Through a Framework



Every one of those ideas exists in Spring, in FastAPI, and in Express. The names and syntax change, the concepts do not. Learn them once in your chosen framework and you can pick up any other on the job in about a week.

You do not need to go super deep here, especially now. Understand how the framework works at a high level, ship a real backend with it, and you are in good shape. The deep knowledge comes naturally once you're employed and using it every day.

Project: Restructure your habit tracker

Take the FastAPI app from part one and structure it like a real project. Split your routes into their own files, add a middleware that logs every incoming request, add proper validation on your request bodies, and add an error handler that returns clean JSON when something goes wrong. Move any config into environment variables.

3. Databases

Right now your habit tracker forgets everything on restart, because it's storing habits in a Python list that vanishes when the program stops. A **database** is a program whose entire job is to store data permanently and let you retrieve it fast. Learning databases well is one of the highest-leverage things you can do as a backend engineer.

Start with **SQL**. Always.

SQL databases store data in tables, like very powerful spreadsheets, and you query them with a language also called SQL. NoSQL databases (which store data in other shapes) are everywhere in industry and you will use them, but a relational SQL database forces you to understand **data modeling**, and data modeling is the actual skill. It is the bread and butter. Once you can model data well in SQL, NoSQL is easy. The reverse is not true.

Use [PostgreSQL](#) (everyone calls it Postgres). It is free, it is everywhere, and it is what a huge chunk of the industry runs on.

Here are the concepts you need to understand:

Concept	What it is
Tables & rows	A table is like a spreadsheet; each row is one record
Columns	The fields each row has (name, email, created date)
Primary keys	The unique ID for each row, so you can point at exactly one
Foreign keys	A column in one table that references a row in another
Indexes	A lookup structure that makes reads fast (and writes slightly slower)
Joins	Combining related data across two or more tables in one query
Constraints	Rules the database enforces for you (not null, must be unique, etc.)

Concept	What it is
Migrations	Versioned changes to your table structure over time
Transactions	A group of operations that all succeed together or all fail together

Every one of these is critical. Indexes and joins are where most beginners are weak, and transactions are where most bugs in real systems hide (imagine moving money between two accounts and the second step fails: a transaction makes sure you never lose money halfway through).

Here is the part that matters most: **learning how to connect and spin up a database is the easy part. Learning how to model data is the hard part, and the important one.** Anyone can run `CREATE TABLE`. The tough part is deciding what tables you need, what columns they have, and how they relate to each other. Get this wrong and you'll one day face the worst pain in backend: a database migration on a live system.

Project: Add Postgres to your habit tracker

Take your habit tracker and put a real Postgres database behind it. Before you write any code, design the tables on paper. You need at least users, habits, and completions. Decide your primary keys, your foreign keys, and how the tables relate (one user has many habits, one habit has many completions).

Then connect it to FastAPI and move your data out of that in-memory list and into the database. Spend more time on the schema than on the connection code.

Follow this guide to get this done: <https://fastapi.tiangolo.com/tutorial/sql-databases/>

Your data is now persistent. But right now anyone who finds your API can read and write anyone's habits. We need to lock it down.

4. Authentication and Authorization

This is the bread and butter of what makes real systems work. Every single application you will ever build needs this. If you understand auth well, you are ahead of a huge number of people who skip past it and build insecure apps.

Two words that sound the same and mean different things:

- **Authentication:** who are you? (logging in, proving your identity)
- **Authorization:** what are you allowed to do? (access control, permissions)

Authentication comes first (prove who you are), then authorization decides what that person can touch. Here's what to learn, in order.


Password storage

Never, ever store passwords as plain text. If you store the literal password and your database leaks, every user is compromised instantly. Instead you store a **hash**: a scrambled, one-way version of the password made by an algorithm like **bcrypt** or **argon2**. One-way means you cannot turn the hash back into the password.

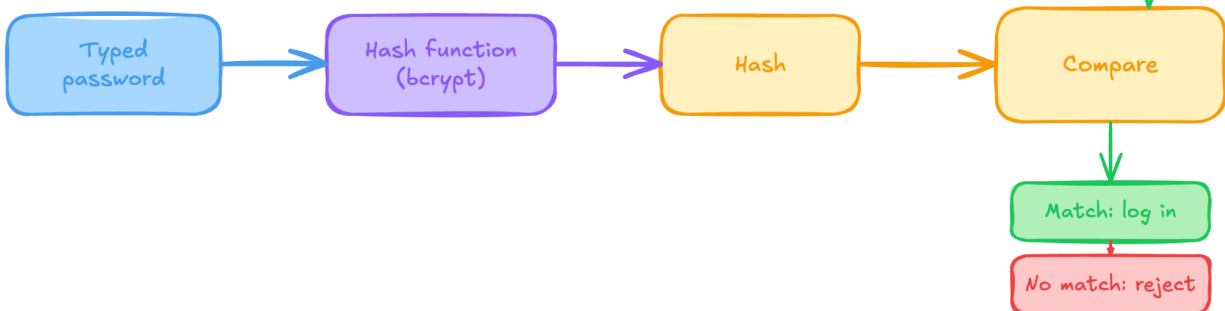
When a user signs up, you hash their password and store only the hash. When they log in, you hash what they just typed and compare it to the stored hash. The database never knows the actual password.

Hashing Passwords on Signup and Login

SIGNUP



LOGIN



Sessions vs JWTs

Once someone logs in, you need a way to remember they're logged in on every following request, since HTTP itself forgets you between requests. Two common approaches:

	Sessions	JWTs
Where the "logged in" state lives	On the server (or a store like Redis)	Inside a token held by the client
Scaling	Needs a shared session store across servers	Stateless, easy to scale
Logging someone out	Easy, just delete the session	Harder, the token is valid until it expires

A **JWT** (JSON Web Token) is a signed string the server hands you when you log in. You send it back on every request to prove who you are. The server can verify it's legit without looking anything up, which is why JWTs scale so well.

Cookies

A **cookie** is a small piece of data the browser stores and sends back automatically, often used to hold a session ID or token. Three settings matter for security, and getting them wrong is a real hole: **HttpOnly** (JavaScript can't read it), **Secure** (only sent over HTTPS), and **SameSite** (limits cross-site sending).

OAuth

OAuth is how "Sign in with Google" or "Sign in with GitHub" works. Instead of storing the user's password yourself, you delegate login to a provider you trust. The high-level flow: your app sends the user to Google, Google confirms who they are and sends back a code, you exchange that code for a token. You never see their Google password, which is safer for everyone.

Access control

Once you know who someone is, you decide what they're allowed to touch. A user should be able to see and edit their own habits, and never anyone else's. This sounds obvious, and forgetting it is one of the most common (and embarrassing) real-world bugs: APIs that happily hand user A all of user B's data because nobody checked.

Project: Add auth to your habit tracker

Add user signup and login. Follow the FastAPI guide for this:

<https://fastapi.tiangolo.com/tutorial/security/first-steps/>

Your API is now quite close to an MVP. Now let's make it hard to break.

5. Input Validation and Error Handling

The first 80% of your project is complete, now you have to get into the nitty gritty of handling edge cases and making stuff work **well**.

Input validation

Never trust input. Ever. Anything coming from a client could be a mistake, garbage, or an outright attack. Most security holes in real systems come down to some form of **injection**, where an attacker sends data that your code accidentally treats as a command instead of plain data.

The classic is **SQL injection**: if you paste user input straight into a database query, an attacker can type input that turns into its own query and reads or wipes your whole database. It happens constantly to careless code.

The fixes are not complicated:

- Validate every incoming request against an expected shape. Your framework helps here, for example **Pydantic** in FastAPI lets you declare exactly what a valid request looks like and auto-rejects the rest.
- Use **parameterized queries** or an **ORM**, never glue user input into a SQL string by hand. This makes the database treat input as data, never as commands.
- Reject anything that doesn't match what you expect, rather than trying to clean it up.

Error handling

Things **will** go wrong. The question is what happens when they do. You want to balance two goals that pull in opposite directions:

1. **Be helpful to yourself.** You want enough detail in your logs to find and fix a problem fast.

2. **Reveal nothing to attackers.** The message that goes back to the caller should be clean and generic. A raw stack trace or database error handed to a user is a gift to a bad actor, it tells them exactly how your system is built.

So the rule is simple: **log the full detail on the server, return a clean message to the client.** The user sees `500 Internal Server Error`, your logs hold the full stack trace. Never the other way around.

Project: Harden your habit tracker

Add strict validation to every endpoint so malformed requests get rejected with a clear `400` and a helpful message. Confirm you're using parameterized queries or an ORM everywhere (if you used an ORM in the database section, you already are). Add a global error handler that logs full detail server-side but returns clean, generic messages to the caller.

Guide to follow for this: <https://fastapi.tiangolo.com/tutorial/handling-errors/>

Your backend is now safer to expose to the world.

6. File Storage and External Services

Real backends rarely live alone. They store files, and they call other companies' services. This section is about doing both safely.

File storage

When users upload things (profile pictures, documents, exports), do not store the files themselves in your database or on your server's hard drive. Both get full, slow, and hard to scale. Instead use an **object store**, a service built specifically to hold files cheaply and serve them fast. You can use whatever you want, but the standard (and what I normally use) is [Amazon S3](#) (or MinIO).

The pattern: the file goes to S3, and you store just the file's **key** (its address in S3) in your database. Your database stays small and fast, S3 holds the heavy stuff.

Calling third-party APIs safely

The moment you call another company's service (a payment processor, an email sender, a weather API), you depend on something you don't control. It will be slow sometimes, and it will fail sometimes. Build like you expect that:

- Always set a **timeout**, so you never wait forever for a service that's hanging
- Add **retries** for temporary failures, but with **backoff** (wait longer between each try) so you don't hammer a struggling service
- Keep API keys in **environment variables**, never written in your code
- **Fail gracefully** when the other service is down, instead of letting it take your whole app down with it

Project: Add file uploads to your habit tracker

Let users upload a profile picture. Store the image in S3 (or run a free local S3-compatible tool like MinIO if you don't want a cloud account yet) and save only the key in Postgres. Then integrate one external API: for example, send a welcome email through a service like Resend or SendGrid when a user signs up, with a timeout and a retry.

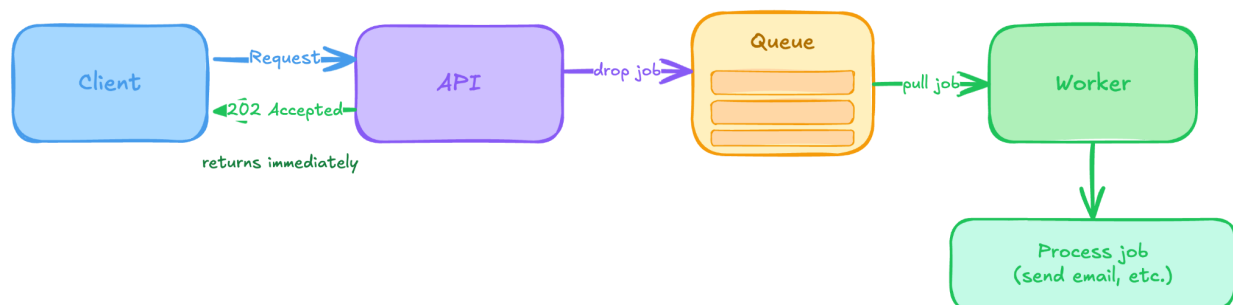
Where to learn it: This step combines two halves. For receiving the file, FastAPI's [Request Files](#) doc shows the `UploadFile` type. For pushing it to storage, you use `boto3`, the [AWS SDK for Python](#), whose `upload_fileobj` call sends the file straight to a bucket. There's no single official doc covering both, so wire them together yourself: receive the `UploadFile`, generate a unique key (use a UUID so two users uploading `photo.jpg` don't collide), upload it, and save just that key to Postgres.

Now your app talks to the outside world. But some of that work is slow, and making users wait for it is a bad experience. Time for queues.

7. Background Jobs and Queues

Some work should not happen while the user waits. Sending an email, resizing an uploaded image, generating a big report. If you do these inside the request, the user stares at a loading spinner and your server is tied up doing slow work instead of serving other people.

The fix is to do the slow work in the background. You put a **job** on a **queue** (a line of work waiting to be done), immediately return a response to the user, and a separate **worker** process pulls jobs off the queue and does them whenever it's ready.



This buys you two big things:

1. **Fast responses.** The user gets an answer immediately, the slow work happens after.
2. **Independent scaling.** If jobs pile up, you just add more workers without touching your API.

Common tools: **Celery** or **RQ** in Python, **BullMQ** in Node, usually with **Redis** or **RabbitMQ** as the queue underneath. As always, nothing is free. With a queue the work is no longer instant, so you need a way for the user to check on it later, usually by giving them a job ID they can use to ask "is it done yet?"

Project: Move email to a background job

Take that welcome email from the last section and move it onto a queue. When a user signs up, drop a job on the queue and return a response immediately. A worker sends the email in the background. Your signup endpoint is now fast no matter how slow the email provider is that day.

Guide to get started with this: <https://testdriven.io/blog/fastapi-and-celery/>

Your backend is fast and resilient. Now let's make it fast even when tons of people are reading the same data.

8. Caching

Caching is one of the simplest, highest-impact ways to speed up a backend and take load off your database. The idea: store frequently-requested data somewhere extremely fast (usually [Redis](#), a database that lives in memory) so you don't have to do the slow work of fetching it every single time.

The catch, and there's always a catch, is **cache invalidation**: knowing when the data in your cache has gone stale and needs to be thrown out or updated. If you cache a user's habit list and they add a habit, the cache is now wrong until you fix it. The simplest starting point is a **TTL (time to live)**: cached data automatically expires after a set time, say 60 seconds. Easy, and good enough for a huge number of cases.

The rule of thumb: only cache things that are read often and change rarely. Caching data that changes every second just serves people stale data and adds complexity for no benefit.

Project: Cache your habit list

Cache a user's habit list in Redis with a short TTL. On a read, check the cache first, and on a miss fall back to Postgres and store the result. Invalidate or update the cache when the user

adds or completes a habit so they never see stale data. Then watch your repeated reads get noticeably faster.

Guide for this:

<https://tutorials.technology/tutorials/redis-python-tutorial-caching-pubsub-2026.html>

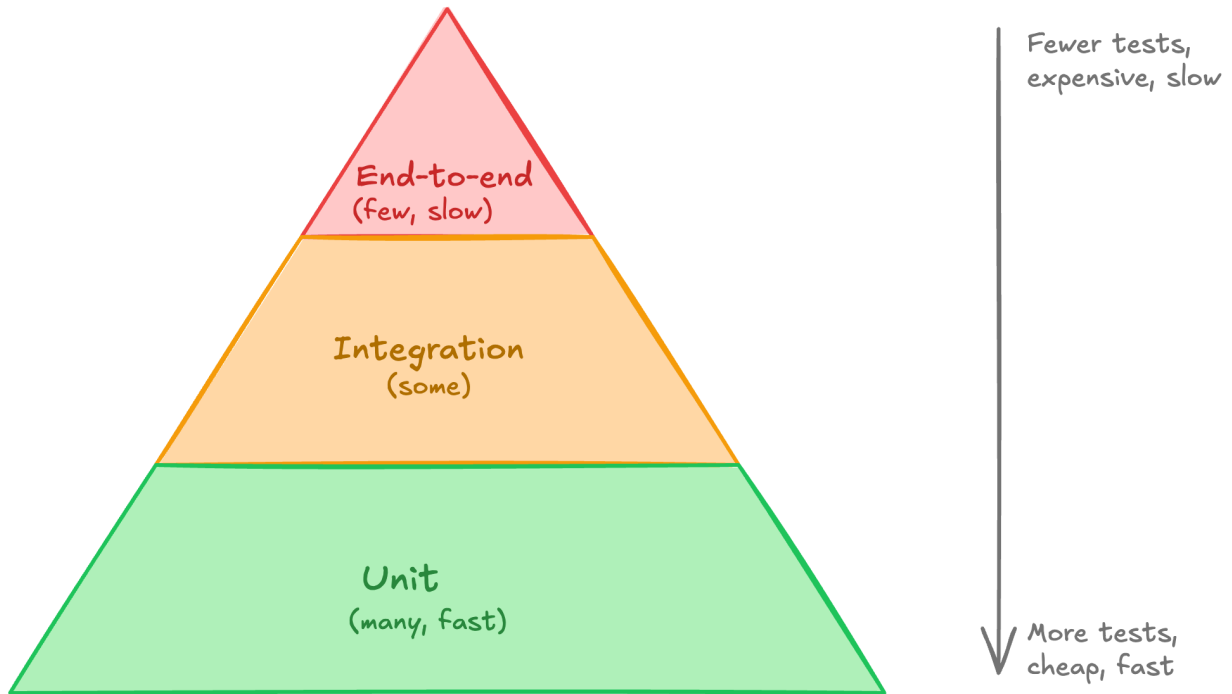
You have a fast, safe, persistent backend. Now let's make sure it actually works, and keeps working as you change it.

9. Testing

Untested code is broken code that just hasn't failed yet. Tests are small programs that automatically check your code does what you think it does. They're what let you change your backend without lying awake wondering what you quietly broke. You know this, but you're probably too lazy to write tests for most of your personal projects. Too bad, you're doing it this time.

There are a few types, and you should understand the differences:

Type	What it tests	Speed
Unit	One small function in isolation	Very fast
Integration	Several pieces working together (your API + the database)	Slower
End-to-end	The whole flow exactly as a real user would hit it	Slowest



You do not need to test every line (chasing 100% coverage is a waste of time for everyone). You need your **critical paths** covered: signup, login, creating a habit, and access control.

Use your language's standard tools: **pytest** for Python, **JUnit** for Java, **Jest** for Node.

Project: Test your habit tracker

Write unit tests for your core logic and integration tests for your main endpoints. At minimum, prove that signup and login work, that a logged-in user can create and complete a habit, and that one user absolutely cannot read or modify another user's habits.

Guide: <https://fastapi.tiangolo.com/tutorial/testing/>

Your backend works and you can prove it. The last step is getting it off your laptop and onto the actual internet.

10. Deployment and Observability

A backend running only on your laptop is a hobby. **Deploying** it (putting it on a server that's always online so anyone can use it) makes it real. **Observability** is how you know it's still alive and healthy once it's out there where you can't watch it directly.

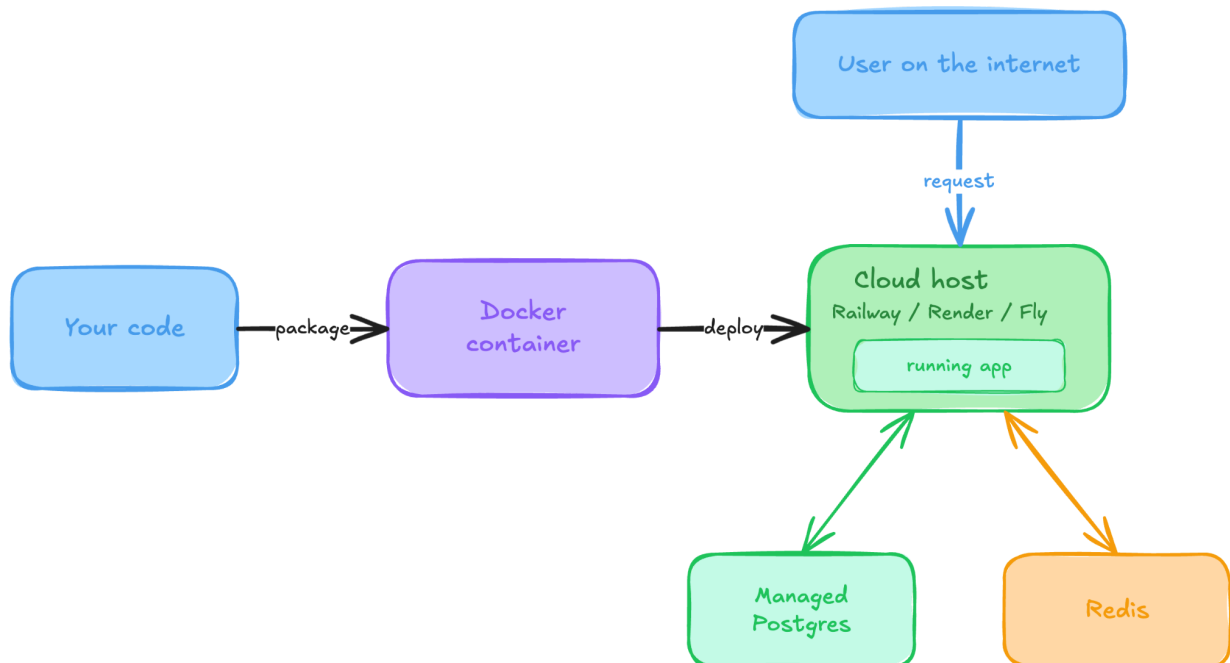
Deployment

The big problem deployment solves: code that runs on your laptop often breaks on a different machine because something is set up differently. The fix is **Docker**, which packages your app and everything it needs into a **container** that runs identically everywhere. You build the container once, and it behaves the same on your laptop and in the cloud.

You do not need Kubernetes or any of the scary stuff on day one. Start simple:

- A beginner-friendly platform like **Railway**, **Render**, or **Fly.io** that takes your container and runs it online for you
- A **managed Postgres** (the platform runs the database for you) instead of operating your own
- **Environment variables** for all your config and secrets, set in the platform's dashboard

A Simple Production Deployment



Observability

Once your app is live, you can't just glance at your terminal to see what it's doing. Observability gives you eyes on it. Three pillars:

- **Logs:** a written record of what happened (you already set these up back in error handling)
- **Metrics:** numbers tracked over time, like requests per second, error rate, and response time
- **Alerts:** automatic notifications that tell you something is wrong before your users do

Structured logs plus one basic alert (for example, "tell me if the 5xx error rate spikes") already put you ahead of a shocking number of real production services.

Project: Deploy your habit tracker

Containerize your habit tracker with Docker. Deploy it to Railway, Render, or Fly.io with a managed Postgres and Redis. Confirm you can hit your live API from your phone, off your home wifi, from anywhere. Text your mom and ask her to give it a shot. Add structured logging and a basic alert for when your error rate spikes.

That's it. You now have a complete, deployed, tested backend with authentication, a real database, file storage, background jobs, caching, and monitoring. That is the shape of a real production service, and you built the entire thing yourself.

What's Next?

If you made it here and actually built the habit tracker step by step instead of just reading, you're no longer learning backend. You're doing backend. The next step is depth and reps.

If you want to understand how these pieces fit together at serious scale (when to choose strong vs eventual consistency, how to design for millions of users instead of one), read my [Ultimate System Design Playbook](#) next. This roadmap teaches you to build most basic service. System design teaches you to build it for the whole world.

And if you want to keep practicing alongside other engineers (and me), check out **The Daily Dev**, my app and community for developers who want to actually get better at this stuff by practicing every day (I just wrapped up today's question). Available on the App Store and the web. See you inside!

<https://thedailydevweb.arjaythede.com/>

