AI VOID FIELD GUIDE

Mastering Trigger.dev: A Zero-to-Advanced Guide for AI Workflows

Embark on a comprehensive journey to master Trigger.dev v4-beta, learning to build, deploy, and manage robust AI agents and automated workflows for modern production systems.

AI VOID

2026-05-20 14:55 IST

Contents

01	Welcome to Trigger.dev v4-beta: The Foundation for Modern Workflows	3
02	Setting Up Your Trigger.dev Environment & First Workflow	14
03	Mastering Basic Workflows: Events, Tasks, and Retries	26
04	Building Robust Workflows: Queues, Scheduling, and Long-Running Processes	41
05	Integrating Trigger.dev: Next.js, TypeScript, and External APIs	55
06	Observability & Debugging: Seeing Your Workflows in Action	70
07	Unleashing AI Agents: Building Smart, Automated Systems	84
08	Human-in-the-Loop & Real-time Updates: Collaborative Workflows	99
09	Advanced Integrations: Understanding MCP & Custom Connectors	115
10	Self-Hosting Trigger.dev: Taking Full Control (Advanced)	130
11	Real-World Project: Building an AI-Powered Customer Support Agent	150

CHAPTER 01

Welcome to Trigger.dev v4-beta: The Foundation for Modern Workflows

Building modern applications, especially those integrating AI, often means dealing with complex, distributed systems. You need to ensure tasks run reliably, recover from failures, and scale gracefully. This is where tools like Trigger.dev shine.

In this introductory chapter, we'll lay the groundwork for mastering Trigger.dev v4-beta. You'll learn what Trigger.dev is, why it's becoming an essential tool for developers, and how to set up your very first project. We'll then walk through creating a simple, durable background job, observing its execution, and understanding the core principles that make Trigger.dev powerful. By the end of this chapter, you'll have a running Trigger.dev project and a foundational understanding of its capabilities.

Why Trigger.dev v4-beta? The Need for Resilient Workflows

Imagine building an AI agent that analyzes customer feedback, generates personalized responses, and then sends those responses via email. What happens if the email service is temporarily down? Or if the AI model takes too long to respond? Without a robust system, your agent might fail, leaving customers in the lurch.

This is the challenge Trigger.dev addresses. It provides a platform for building **durable, reliable, and observable server-side workflows** that can:

- **Handle failures gracefully:** Automatically retry tasks, back off on errors, and ensure eventual success.
- **Manage long-running operations:** Keep track of state for workflows that might take minutes, hours, or even days to complete.
- **Orchestrate complex sequences:** Define steps that run in order, in parallel, or conditionally, making complex logic manageable.
- **Integrate with anything:** Connect to APIs, databases, message queues, and AI models seamlessly.

Trigger.dev v4-beta represents the next generation of this platform. While v3 is the current stable release, v4-beta introduces significant improvements and is expected to reach General Availability (GA) around May/June 2026. By starting with v4-beta now, you'll be ahead of the curve, learning the most modern patterns and features.

Key Concepts Explained

Let's quickly define some foundational concepts that Trigger.dev builds upon:

- **Background Jobs:** Tasks that run asynchronously, often in response to an event, without blocking the main application thread. Think sending an email or processing an image.
- **Durable Execution:** The ability for a workflow to pause, persist its state, and resume later from where it left off, even if the server restarts. This is critical for long-running processes.
- **Queues:** A mechanism to hold tasks before they are processed, ensuring that work is distributed and not lost.
- **Retries:** Automatically re-attempting failed operations, often with exponential backoff, to overcome transient issues.
- **Scheduling:** Running jobs at specific times or intervals (e.g., daily reports, hourly data syncs).
- **Observability:** The ability to understand the internal state of your workflows, including their progress, logs, and any errors. This is crucial for debugging and monitoring.
- **AI Agents:** Complex workflows that involve multiple steps, often interacting with AI models, external tools, and potentially human input, to achieve a specific goal.

MCP Integration: Trigger.dev is designed to be highly interoperable. While "MCP" isn't a universally defined acronym in this context, within Trigger.dev's ecosystem, it broadly refers to **Multi-Cloud Platform** or **Managed Cloud Platform** integration. This emphasizes Trigger.dev's capability to seamlessly connect your workflows to various external services, cloud providers, and managed platforms, including diverse AI models, databases, and APIs. This flexibility allows you to build comprehensive agentic systems that leverage the best tools from across the cloud landscape.

Setting Up Your First Trigger.dev Project (v4-beta)

Ready to get your hands dirty? Let's initialize a new project.

Prerequisites:

Before we begin, ensure you have:

- **Node.js (v18 or higher):** As of 2026-05-20, Node.js v20 is the current LTS. You can download it from nodejs.org. Using a version manager like `nvm` (Node Version Manager) is highly recommended for easily switching between Node.js versions.
- **npm** (comes with Node.js) or **Yarn**.

Step 1: Initialize Your Project

Open your terminal and navigate to the directory where you want to create your project. Then, run the following command to initialize a new Trigger.dev project using the v4-beta version:

```
npx trigger.dev@v4-beta init
```

What's happening here?

- `npx`: Executes a Node.js package without explicitly installing it globally.
- `trigger.dev@v4-beta`: Specifies that we want to use the `trigger.dev` package and specifically the `v4-beta` version. This is crucial for accessing the latest features we'll be exploring.
- `init`: This is the command to scaffold a new Trigger.dev project.

The `init` command will ask you a few questions:

1. **Project Name:** You can press Enter to accept the default or provide your own (e.g., `my-first-trigger-project`).
2. **API Key:** It will prompt you to create an API key on `trigger.dev`. Follow the link provided in your terminal, sign up/log in, and create a new project. Once created, copy the `TRIGGER_API_KEY` and `TRIGGER_PROJECT_ID` and paste them back into your terminal.
3. **Framework:** Choose `Next.js` for now. Trigger.dev integrates well with various frameworks, but Next.js is a common choice for modern web applications and provides a good starting point.
4. **Language:** Choose `TypeScript`. This guide will primarily use TypeScript for its type safety and developer experience benefits.

Once completed, you'll see a new directory created with your project name.

Step 2: Explore the Project Structure

Navigate into your new project directory:

```
cd my-first-trigger-project # Or whatever you named your project
```

Take a moment to look at the generated files:

- `package.json`: Standard Node.js project configuration.
- `.env`: Contains your `TRIGGER_API_KEY` and `TRIGGER_PROJECT_ID`. **Never commit this file to version control!**
- `trigger.config.ts`: This is the main configuration file for your Trigger.dev client.
- `src/trigger.ts`: Where your Trigger.dev client instance is created and exported.
- `src/jobs/`: This directory is where you'll define your individual jobs.

Step 3: Run the Development Server

To start your Trigger.dev development server and connect it to the Trigger.dev cloud dashboard, run:

```
npm run dev
```

You should see output indicating that your Next.js application is running and that the Trigger.dev client is connected.

What does this do?

- This command starts your local development server.
- Your local Trigger.dev client (running within your Next.js app) establishes a connection to the Trigger.dev cloud service using your API key.
- This connection allows your local jobs to be registered with the Trigger.dev dashboard and enables real-time logging and execution monitoring.

Leave this terminal window open, as your Trigger.dev client needs to be running to execute jobs.

Your First Workflow: A Simple Background Job

Now that our project is set up, let's create a basic job. In Trigger.dev, a "Job" is a function that you define to perform a specific task. These jobs are durable, meaning Trigger.dev ensures they run to completion, even if your local server temporarily disconnects or restarts.

Step 1: Create a New Job File

Inside your `src/jobs/` directory, create a new file named `helloWorld.ts`:

```
// src/jobs/helloWorld.ts
import { client } from "@/trigger";
import { eventTrigger } from "@trigger.dev/sdk";

// Define your first job!
// This job will simply log a message when triggered.
client.defineJob({
  // A unique ID for your job.
  // This is used to identify the job in the Trigger.dev dashboard.
  id: "hello-world-job",
  // A display name for your job, more human-readable.
  name: "Hello World Job",
  // The version of your job definition.
  // Increment this when you make significant changes.
  version: "0.1.0",
  // The trigger that starts this job.
  // `eventTrigger` means it starts when a specific event is received.
  trigger: eventTrigger({
    name: "hello.world", // The name of the event that will trigger this job
  }),
  // The actual function that runs when the job is triggered.
  async run(payload, io, ctx) {
    // `payload` contains the data sent with the event.
    // `io` is the I/O client for interacting with external services and
logging.
    // `ctx` provides context about the job run.

    // Use `io.logger` for logging within your jobs.
    // This makes logs visible in the Trigger.dev dashboard.
    io.logger.info("Hello from Trigger.dev!", {
      timestamp: new Date().toISOString(),
      receivedPayload: payload,
    });

    // You can also return data from your job.
    return {
      message: "Job completed successfully",
      payloadProcessed: payload,
    };
  },
});
```

Explanation of the code:

- `import { client } from "@/trigger";`: Imports your Trigger.dev client instance. This client is how you register jobs with the Trigger.dev platform.
- `import { eventTrigger } from "@trigger.dev/sdk";`: Imports the `eventTrigger` helper, which tells Trigger.dev that this job should run whenever a specific event is sent to it.
- `client.defineJob({...});`: This is the core function for defining a job.
    - `id`: A unique, machine-readable identifier for your job.
    - `name`: A human-readable name displayed in the dashboard.
    - `version`: Helps you track changes to your job definitions.
    - `trigger`: Defines how the job starts. Here, `eventTrigger({ name: "hello.world" })` means it will activate when an event named `hello.world` is received.
    - `run(payload, io, ctx)`: This is the main function where your job's logic resides.
        - `payload`: Any data sent with the triggering event.
        - `io`: An object providing utilities for logging (`io.logger`), interacting with external services, and managing state. This is key to making your jobs observable and durable.
        - `ctx`: Provides contextual information about the job run itself.
- `io.logger.info(...)`: This is how you log messages within a Trigger.dev job. These logs are captured by Trigger.dev and are visible in the dashboard, making debugging much easier than traditional `console.log`.

Step 2: Trigger the Job from the Dashboard

1. Make sure your local development server is still running (`npm run dev`).
2. Open your web browser and navigate to the Trigger.dev dashboard (the URL provided during setup, typically `<https://cloud.trigger.dev/>`).
3. Log in and select your project.
4. You should see your "Hello World Job" listed under the "Jobs" section. If not, it might take a moment to register, or there might be an issue with your `npm run dev` output (check for errors).
5. Click on your "Hello World Job".
6. On the job details page, you'll find a "Test" tab or a "Run Job" button. Click it.

7. You'll be prompted to enter a "Payload". This is the data that will be passed to your job's `run` function. Enter some JSON, for example:

```
{
  "who": "World",
  "message": "Greetings from the dashboard!"
}
```

1. Click "Run Job".

Step 3: Observe the Execution

After you run the job from the dashboard:

1. Navigate to the "Runs" tab for your job in the Trigger.dev dashboard.
2. You should see a new job run listed with a "Success" status.
3. Click on the run to view its details. Here, you'll see:
    - The `payload` you sent.
    - The `output` returned by your job.
    - Crucially, the `io.logger.info` messages you wrote, along with their timestamps. This is your observability in action!

Congratulations! You've successfully defined, triggered, and observed your first durable background job with Trigger.dev v4-beta.

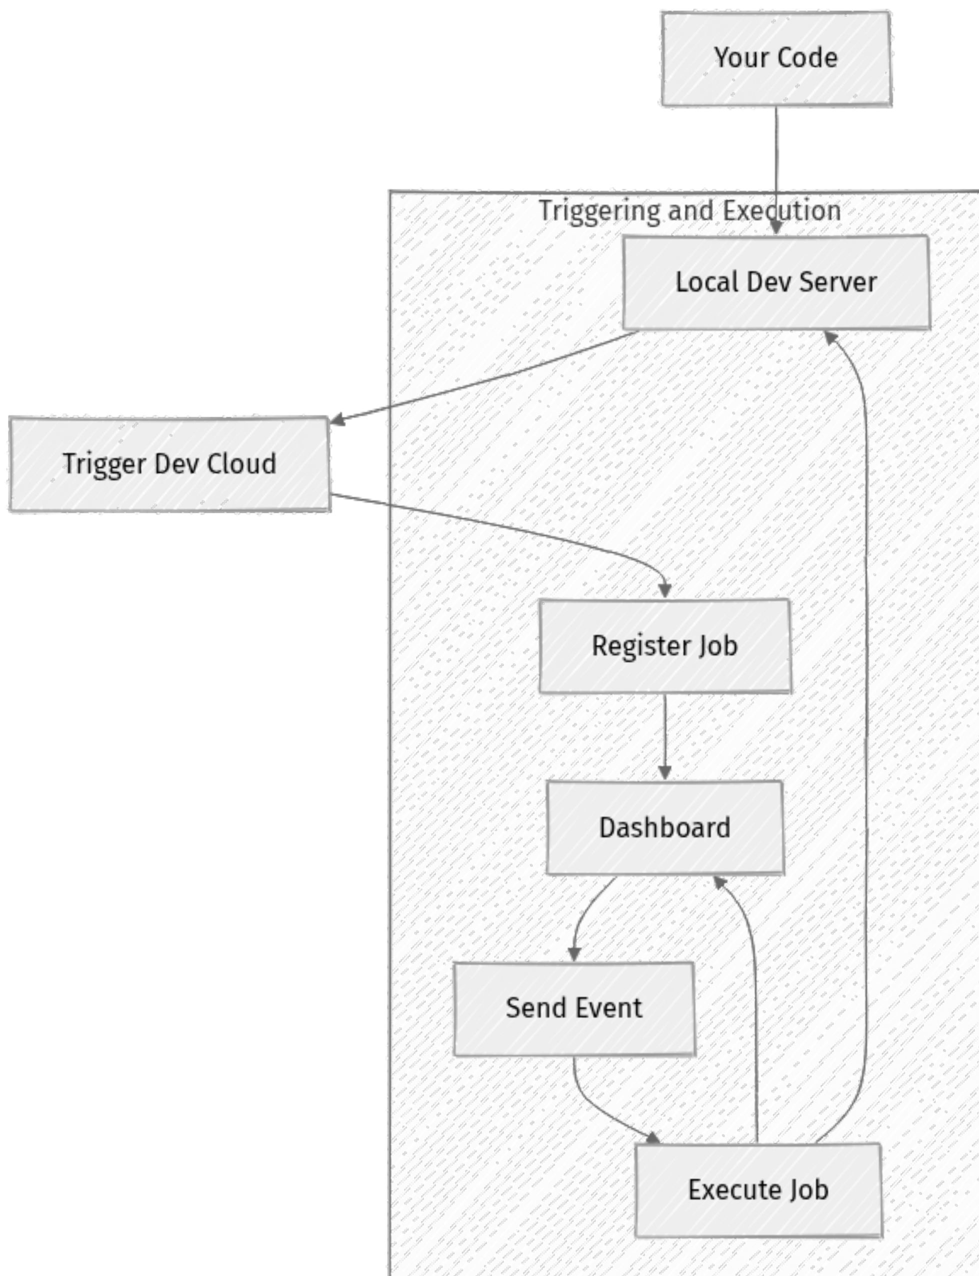

Mini-Challenge: Personalize Your Greeting

Let's make our `hello-world-job` a bit more dynamic.

Challenge: Modify the `helloWorld.ts` job to greet the `who` property from the `payload` you send, instead of just "Hello from Trigger.dev!". If `who` is not provided in the payload, it should default to "Stranger".

Hint:

- Access properties from the `payload` object directly (e.g., `payload.who`).
- Use a default value if `payload.who` is `undefined` or `null` (e.g., `payload.who || "Stranger"`).

What to Observe/Learn:

- How to access and use input data (`payload`) within your job logic.
- How changes to your job definition are picked up by the running `npm run dev` process (usually automatically, but sometimes a restart is needed).
- How different payloads affect the job's output and logs in the dashboard.

After modifying your code, save the file. The `npm run dev` server should automatically reload. Then, go back to the Trigger.dev dashboard, trigger the job again with a payload like `{"who": "Alice"}`, and then again with an empty payload `{}`. Check the logs each time!

Common Pitfalls & Troubleshooting

Even with simple setups, you might encounter issues. Here are a few common ones for beginners:

1. **Wrong Trigger.dev Version:**
    - **Pitfall:** Accidentally using `npx trigger.dev init` instead of `npx trigger.dev@v4-beta init`. This will set up a v3 project, and the code examples in this guide might not work as expected.
    - **Troubleshooting:** Always ensure you specify `@v4-beta` during initialization. If you've started with v3, it's best to create a new project with v4-beta for this guide.
2. **Missing or Incorrect Environment Variables:**
    - **Pitfall:** Your `.env` file doesn't have `TRIGGER_API_KEY` and `TRIGGER_PROJECT_ID` set correctly, or you committed it to Git and deployed without it.
    - **Troubleshooting:** Verify these variables are present and correct in your local `.env` file. If running in a deployed environment, ensure they are configured as environment variables there. Your `npm run dev` output will usually show an error if the client can't connect.
3. **Local Dev Server Not Running:**
    - **Pitfall:** You closed the terminal where `npm run dev` was running, or it crashed. Your jobs won't be registered, and events won't be processed.
    - **Troubleshooting:** Always ensure `npm run dev` is active in a terminal window while you are developing and testing jobs locally.

4. **Job Not Showing in Dashboard:**
    - **Pitfall:** You defined a job but it's not appearing in the "Jobs" list on the Trigger.dev dashboard.
    - **Troubleshooting:**
        - Check your terminal for any errors from `npm run dev`.
        - Ensure the job file is correctly imported or discovered by your `src/trigger.ts` (the default `init` setup usually handles this).
        - Sometimes, a full restart of `npm run dev` is needed after adding a new job file.
        - Verify your `TRIGGER_PROJECT_ID` matches the project you're viewing in the dashboard.

Summary

In this foundational chapter, you've taken your first steps into the world of Trigger.dev v4-beta. We covered:

- The importance of **durable execution** and **resilient workflows** in modern application development, especially with AI.
- An introduction to Trigger.dev's core concepts, including **background jobs, queues, retries, and observability**.
- The step-by-step process to **initialize a new Trigger.dev v4-beta project** using `npx trigger.dev@v4-beta init`.
- How to **define your first simple background job** and trigger it from the Trigger.dev dashboard.
- Key methods for **observing job execution** and logs.
- Common **pitfalls and troubleshooting** tips for initial setup.

You've built a solid foundation. In the next chapter, we'll dive deeper into different types of triggers, explore how to build more complex workflows, and begin to understand how Trigger.dev handles long-running processes with durable execution. Get ready to unlock even more power!

References

- Trigger.dev GitHub Repository: https://github.com/triggerdotdev/trigger.dev
- Trigger.dev Documentation: https://trigger.dev/docs
- Node.js Official Website: https://nodejs.org/
- TypeScript Handbook: https://www.typescriptlang.org/docs/handbook/

This page is AI-assisted and reviewed. It references official documentation and recognized resources where relevant.

CHAPTER 02

Setting Up Your Trigger.dev Environment & First Workflow

Welcome to Chapter 2! In the previous chapter, we explored the "why" behind Trigger.dev, understanding its role in building robust, fault-tolerant AI agents and automated workflows. Now, it's time to roll up our sleeves and dive into the "how."

This chapter will guide you through setting up your local development environment for Trigger.dev v4-beta and creating your very first job. By the end, you'll have a running Trigger.dev project, a basic understanding of its core components, and the satisfaction of seeing your first durable workflow execute. This hands-on experience is crucial for building confidence and understanding how Trigger.dev fits into your development stack.

Preparing Your Development Environment

Before we jump into Trigger.dev, let's ensure your local machine has the necessary tools. Think of this as preparing your workshop before starting a new project.

Prerequisites: Node.js and npm

Trigger.dev projects typically run on Node.js and leverage `npm` (Node Package Manager) or `Yarn` for dependency management.

1. **Node.js (LTS Version):** Trigger.dev requires Node.js version 18 or higher. The current Long Term Support (LTS) release is always recommended for stability and ongoing support.
    - **Why it matters:** Newer Node.js versions bring performance improvements, security fixes, and modern JavaScript features that Trigger.dev utilizes. Using an older version might lead to compatibility issues or missing features.
    - **Check your version:** Open your terminal or command prompt and run:

```
node -v
```

```
- If your version is older than 18, you'll need to upgrade. We recommend using a Node Version Manager like `nvm` for easy switching between Node.js versions.
    - **Official Node.js Downloads:** [<https://nodejs.org/en/download]>(https://nodejs.org/en/download)
    - **nvm (Node Version Manager):** [<https://github.com/nvm-sh/nvm]>(https://github.com/nvm-sh/nvm)
```

1. **npm (or Yarn):** `npm` is usually installed automatically when you install Node.js.
    - **Check your version:**

```
npm -v
```

```
- If you prefer Yarn, ensure it's installed globally:
```

```
yarn -v
```

```
- ⚡ Quick Note: While `npm` is the default, Trigger.dev supports `yarn` as well. The initialization process will adapt to your preference if you have Yarn installed.
```

Core Concepts: Building Blocks of a Trigger.dev Workflow

Before we write any code, let's establish a mental model for what we're about to build. Understanding these core concepts will make the implementation steps much clearer.

What is a Trigger.dev Job?

At its heart, a Trigger.dev job is a durable, fault-tolerant function that executes in response to an event.

- **What it is:** A unit of work defined in your code that Trigger.dev manages. Think of it as a specific task, like "send a welcome email" or "process a payment."
- **Why it exists:** To perform tasks reliably, even if your server crashes, the network goes down, or an external API is temporarily unavailable. Trigger.dev ensures the job either completes successfully or is retried until it does, providing robust operations in unpredictable environments.

- **How it works:** You define a job, specifying what event triggers it and what code it should run. Trigger.dev then orchestrates its execution, including retries, delays, and state management, ensuring its eventual completion.

Triggers and Events

Every Trigger.dev job starts with a trigger. This is how your job knows when to spring into action.

- **Trigger:** The condition that causes a job to start. This could be a scheduled interval (e.g., "every hour"), an incoming HTTP webhook event (e.g., "when a user signs up"), or a custom event you send programmatically.
- **Event:** The actual data payload that initiates a job. When a trigger fires, it often carries an event with relevant data for the job to process. For example, a `user.signed.up` event might carry the user's ID and email address, which the job then uses to personalize a welcome message.

The Trigger.dev CLI

The Trigger.dev Command Line Interface (CLI) is your primary tool for initializing projects and interacting with the Trigger.dev platform.

- **What it is:** A command-line utility that helps you set up and manage your Trigger.dev projects.
- **Why it exists:** To streamline the initial setup, ensuring you have the correct project structure and dependencies, and to connect your local development environment to the Trigger.dev cloud service. It automates much of the boilerplate.
- **How it works:** We'll use `npx trigger.dev@v4-beta init` to scaffold a new project, which handles installing the necessary SDKs and creating boilerplate configuration files. This command allows you to use the CLI without a global installation.

Step-by-Step Implementation: Your First Workflow

Let's get hands-on and create a simple "Hello World" job. This will demonstrate the full lifecycle from setup to execution.

Step 1: Create Your Project Directory

First, create a new directory for your Trigger.dev project and navigate into it. This keeps your project organized and separate from other codebases.

```
mkdir my-first-trigger-project
cd my-first-trigger-project
```

Step 2: Initialize Trigger.dev v4-beta

Now, we'll use the Trigger.dev CLI to set up our project. We'll specifically target the `v4-beta` version, as it's the latest cutting-edge release, with General Availability (GA) expected around May/June 2026. This ensures you're working with the most modern features.

```
npx trigger.dev@v4-beta init
```

- **What's happening here?**
  - `npx`: This command executes a Node.js package without requiring you to globally install it first. It's great for one-off commands or trying out new tools.
  - `trigger.dev@v4-beta`: Specifies the Trigger.dev package and explicitly requests its `v4-beta` version.
  - `init`: This is the command to initialize a new Trigger.dev project, setting up all the necessary files and configurations.

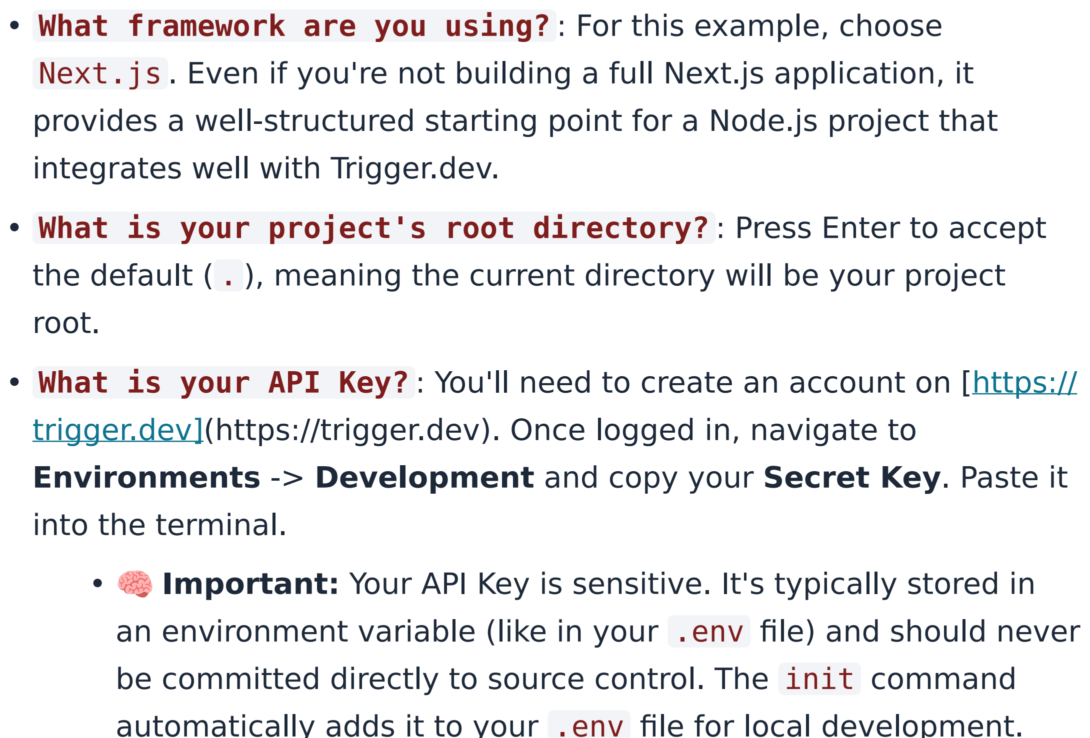

- **Interactive Setup:** The CLI will ask you a few questions to configure your project:
  - **`What framework are you using?`**: For this example, choose `Next.js`. Even if you're not building a full Next.js application, it provides a well-structured starting point for a Node.js project that integrates well with Trigger.dev.
  - **`What is your project's root directory?`**: Press Enter to accept the default (`.`), meaning the current directory will be your project root.
  - **`What is your API Key?`**: You'll need to create an account on https://trigger.dev. Once logged in, navigate to **Environments** -> **Development** and copy your **Secret Key**. Paste it into the terminal.
    - 🧠 **Important:** Your API Key is sensitive. It's typically stored in an environment variable (like in your `.env` file) and should never be committed directly to source control. The `init` command automatically adds it to your `.env` file for local development.

- **Output:** After answering, the CLI will install dependencies and create several files and folders. You should see output similar to:

```
    ✔ What framework are you using? › Next.js
    ✔ What is your project's root directory? › .
    ✔ What is your API Key? › sk_...

    Installing dependencies...
    ... (npm/yarn install output) ...

    🎉 Trigger.dev project initialized!

    Next steps:
    1. Start your dev server: npm run dev
    2. Go to your Trigger.dev dashboard to see your project: https://
app.trigger.dev/orgs/...
```

Step 3: Understanding the Generated Project Structure

Let's quickly look at the key files and folders the `init` command created. This gives you a map of your new Trigger.dev project.

- `package.json`: Your project's manifest file, listing scripts (like `dev`) and dependencies (e.g., `@trigger.dev/sdk`, `@trigger.dev/nextjs`).
- `.env`: Contains your `TRIGGER_SECRET_KEY` environment variable. This file is excluded from version control by default.
- `trigger.config.ts`: The main configuration file for your Trigger.dev project. This is where you configure things like your API key, base URL, and integrations.
- `src/trigger/client.ts`: Exports your initialized Trigger.dev client instance. This client is the bridge between your code and the Trigger.dev platform, used to define and register your jobs.
- `src/trigger/index.ts`: This is where you'll define your Trigger.dev jobs. The `init` command might have added a sample job here, which we will modify or replace.

Step 4: Creating Your First "Hello World" Job

Now, let's open `src/trigger/index.ts` in your code editor and add a very simple job. If there's already a sample job, you can either modify it or add a new one alongside it.

Open `src/trigger/index.ts`:

```
// src/trigger/index.ts
import { trigger } from "./client";
```

```
import { eventTrigger } from "@trigger.dev/sdk";

// Define your first Trigger.dev job
trigger.defineJob({
  // A unique identifier for your job.
  // This is how Trigger.dev tracks and manages it across runs.
  id: "hello-world",
  // A human-readable name for your job, visible in the Trigger.dev dashboard.
  name: "Hello World Job",
  // Versioning helps manage changes to your job definitions over time.
  // It allows for graceful updates without breaking in-flight jobs.
  version: "1.0.0",
  // The 'on' property defines what event will trigger this job.
  // Here, we're using eventTrigger to listen for a custom named event.
  on: eventTrigger({
    // The name of the event this job listens for.
    // Events are typically namespaced for clarity (e.g., "user.created",
"order.processed").
    name: "hello.world",
  }),
  // The 'run' function contains the actual logic of your job.
  // It's an async function because jobs often involve asynchronous operations
  // like API calls, database interactions, or delays.
  run: async (payload, io, ctx) => {
    // 'io.logger' is a special logger provided by Trigger.dev.
    // Messages logged here will appear in your Trigger.dev dashboard for this
specific job run,
    // making debugging distributed workflows much easier than standard
console.log.
    io.logger.info("Hello, Trigger.dev!");
    io.logger.info("Received payload:", payload);

    // You can return any data from your job. This data will be stored as the
job's output
    // and is visible in the Trigger.dev dashboard.
    return { message: "Success! Hello from your first Trigger.dev job!" };
  },
});

// You can define more jobs here by calling trigger.defineJob again...
// For example:
// trigger.defineJob({
//   id: "another-job",
//   name: "Another Example Job",
//   version: "1.0.0",
//   on: eventTrigger({ name: "another.event" }),
//   run: async (payload, io, ctx) => {
//     io.logger.info("Another job ran!");
//     return { status: "completed" };
//   },
// });
```

- **`import { trigger } from "./client";`**: This line imports the `trigger` client instance that was initialized in `src/trigger/client.ts`. This client is your gateway to defining and interacting with the Trigger.dev platform.

- **`import { eventTrigger } from "@trigger.dev/sdk";`**: This imports the `eventTrigger` helper. It's a convenient way to specify that your job should activate when a particular named event is received.
- **`trigger.defineJob({...})`**: This is the core function for defining any Trigger.dev job. It takes a configuration object:
    - **`id`**: A unique string identifier. This is crucial for Trigger.dev to track, manage, and display your job in the dashboard.
    - **`name`**: A human-readable name for the job, which helps you identify it easily in the Trigger.dev dashboard.
    - **`version`**: A semantic version string (e.g., "1.0.0") for your job's logic. This is important for durable execution, allowing Trigger.dev to manage different versions of your job gracefully.
    - **`on: eventTrigger({ name: "hello.world" })`**: This tells Trigger.dev to run this job whenever an event with the name `hello.world` is received.
    - **`run: async (payload, io, ctx) => {...}`**: This is the heart of your job – the function containing the actual business logic.
        - `payload`: An object containing the data sent with the incoming event.
        - `io`: An object providing I/O operations managed by Trigger.dev, such as durable logging (`io.logger`), performing retriable API calls (`io.runTask`), and interacting with integrations. This is key for fault tolerance.
        - `ctx`: Provides contextual information about the current job run, such as `ctx.id` (the unique ID for this specific run).
    - **`io.logger.info(...)`**: Unlike `console.log`, `io.logger.info` is Trigger.dev's durable logger. Logs here are captured, associated with the specific job run, and displayed in your Trigger.dev dashboard, making debugging distributed and long-running workflows significantly easier.

Step 5: Running Your Development Server

Now that we've defined our job, let's start the development server. This will register your `hello-world` job with the Trigger.dev platform and allow it to listen for incoming events.

In your terminal, from your project's root directory:

```
npm run dev
```

- **What's happening?** This script (defined in your `package.json`) typically starts your Node.js application. The Trigger.dev SDK within your application will connect to the Trigger.dev cloud service using your `TRIGGER_SECRET_KEY` (from the `.env` file) and register all the jobs you've defined. This connection ensures that Trigger.dev knows about your jobs and can orchestrate their execution.
- **Output:** You should see output indicating that your server is running and that Trigger.dev has connected and registered your jobs. Look for messages like "Trigger.dev client connected" and "Registered job: hello-world". Keep this terminal window open; it's your local job worker.

Step 6: Triggering Your Job

With your development server running, your `hello-world` job is now active and waiting for a `hello.world` event. Let's trigger it from the Trigger.dev dashboard!

1. **Open the Trigger.dev Dashboard:** Go to https://app.trigger.dev and log in.
2. **Navigate to your Project:** You should see your project listed on the dashboard. Click on it to enter your project's overview.
3. **Go to the "Events" Tab:** In the left sidebar of your project dashboard, find and click on "Events".
4. **Send an Event:** Click the "Send Event" button (or similar) to manually send an event.
    - **`Event Name`**: Type `hello.world`. This must exactly match the `name` property you defined in your `eventTrigger` in `src/trigger/index.ts`.
    - **`Payload`**: You can send some JSON data. For now, let's keep it simple:

```
{
  "greeting": "world"
}
```

```
- Click "Send Event".
```

- **Observation:**
    - Immediately, your running `npm run dev` terminal should show your `io.logger.info` messages: "Hello, Trigger.dev!" and "Received payload: { greeting: 'world' }".
    - In the Trigger.dev dashboard, go to the "Runs" tab. You should see a new entry for your "Hello World Job" with a "Succeeded" status. Click on this run to see the detailed logs, including your `io.logger.info` messages and the job's final output.

Congratulations! You've successfully set up your Trigger.dev environment, defined a durable job, and seen it execute. This is the foundation for all the powerful workflows you'll build.

Mini-Challenge: Personalizing Your Greeting

Let's make our "Hello World" job a little more interactive by using the data from the event payload.

- **Challenge:** Modify the `hello-world` job to accept a `name` property in its event payload. Instead of just logging "Hello, Trigger.dev!", it should log a personalized greeting like "Hello, [name] from Trigger.dev!".
    - **Steps:**
        1. Open and edit `src/trigger/index.ts`.
        2. Access the `name` property from the `payload` object within the `run` function.
        3. Update the `io.logger.info` message to include the personalized name.
        4. **Restart your `npm run dev` server** (Ctrl+C to stop, then `npm run dev` again) to register the updated job definition with Trigger.dev. This step is crucial for changes to take effect.
        5. Trigger the job again from the Trigger.dev dashboard's "Events" tab, but this time, send a payload like:

```
          {
            "name": "Alice"
```

```
        }
```

- **Hint:** The `payload` object is just a standard JavaScript object. You can access its properties using dot notation (e.g., `payload.name`). Remember to handle cases where `payload.name` might not be provided (e.g., use a default greeting).
- **What to observe/learn:** This challenge reinforces how to pass dynamic data to your jobs and access it within the `run` function, which is fundamental for building useful and context-aware workflows. Observe the new personalized greeting in both your terminal and the Trigger.dev dashboard logs for the specific job run.

Common Pitfalls & Troubleshooting

Even simple setups can hit snags. Here are a few common issues you might encounter during initial setup and how to resolve them.

- **`npm run dev` Not Connecting to Trigger.dev:**
    - **Symptom:** Your terminal shows "Trigger.dev client not connected", "Failed to connect to Trigger.dev", or similar errors, and your jobs don't appear in the dashboard.
    - **Cause:** The `TRIGGER_SECRET_KEY` in your `.env` file is incorrect, has leading/trailing spaces, or is missing. There might also be network issues preventing your local machine from reaching `api.trigger.dev`, or a firewall blocking outbound connections.
    - **Solution:** Double-check your `TRIGGER_SECRET_KEY` against the key displayed in your Trigger.dev dashboard (under Environments -> Development). Ensure there are no extra characters. Verify your internet connection. If on a corporate network, check for proxy settings or firewall rules that might restrict outbound traffic.

- **Job Not Showing in Dashboard / Not Triggering:**
    - **Symptom:** You've run `npm run dev`, and your terminal shows the client connected, but your job doesn't appear in the dashboard's "Jobs" list, or sending an event doesn't cause it to run.
    - **Cause:** The `npm run dev` process isn't actively running, or you haven't restarted it after making changes to your job definition. The `id` of your job in `trigger.defineJob` might not match what you expect, or, most commonly, the `eventTrigger` name (`hello.world`) doesn't exactly match the event name you're sending from the dashboard (it's case-sensitive!).
    - **Solution:** Ensure your `npm run dev` terminal is active and shows "Registered job: [your-job-id]". Always restart your dev server (`Ctrl+C` then `npm run dev`) after modifying job definitions. Carefully compare the `id` and `eventTrigger.name` in your code with how you're interacting with the dashboard.
- **Node.js Version Errors:**
    - **Symptom:** `npm install` or `npm run dev` fails with errors related to Node.js version, such as "Engine 'node' is incompatible with this module."
    - **Cause:** Your installed Node.js version is older than the minimum required (Node.js 18+).
    - **Solution:** Upgrade Node.js to version 18 or higher using `nvm` (recommended for managing multiple Node.js versions) or the official installer from nodejs.org.

Summary

In this chapter, you've taken the crucial first steps in your Trigger.dev journey! You've moved from understanding the concepts to hands-on implementation, laying a solid foundation for more complex workflows.

Here's a quick recap of what we covered:

- **Environment Setup:** We ensured your local development environment was ready, verifying Node.js (v18+) and npm/Yarn installations.
- **Project Initialization:** You used the powerful `npx trigger.dev@v4-beta init` command to scaffold a new Trigger.dev project, setting up essential configuration and client files.

- **Core Concepts:** We built a mental model for Trigger.dev by understanding what a job is, how triggers and events work, and the role of the CLI.
- **First Job Creation:** You defined a simple "Hello World" job in `src/trigger/index.ts`, learning about `defineJob`, `eventTrigger`, `payload`, and the durable `io.logger`.
- **Execution & Observation:** You learned how to start your development server with `npm run dev` and trigger your job manually from the Trigger.dev dashboard, carefully observing its execution and logs.
- **Practical Application:** You successfully completed a mini-challenge to personalize your job's output, reinforcing how to pass and utilize dynamic data.
- **Troubleshooting:** We addressed common setup pitfalls like connection issues, job registration problems, and Node.js version mismatches.

You now have a foundational understanding of how to get a Trigger.dev project up and running, and how to define and execute your first durable background job. In the next chapter, we'll dive deeper into the `io` object, exploring how to perform external API calls reliably and handle retries, which is essential for building robust AI agents and complex, fault-tolerant workflows.

References

- Trigger.dev GitHub Repository: https://github.com/triggerdotdev/trigger.dev
- Trigger.dev Documentation: https://trigger.dev
- Node.js Official Website: https://nodejs.org
- nvm (Node Version Manager): https://github.com/nvm-sh/nvm

This page is AI-assisted and reviewed. It references official documentation and recognized resources where relevant.

CHAPTER 03

Mastering Basic Workflows: Events, Tasks, and Retries

Welcome back! In the previous chapter, we successfully set up our Trigger.dev project, getting ready to build powerful automated systems. Now, it's time to dive into the fundamental building blocks that make Trigger.dev workflows so resilient and effective: **Events**, **Tasks**, and **Retries**. These three concepts are the bedrock for creating robust, automated workflows and AI agents that gracefully handle the complexities and inevitable failures of real-world production environments.

This chapter will guide you through understanding what events are, how tasks execute reliably, and how Trigger.dev automatically handles failures through intelligent retries. By the end, you'll be able to create your first resilient workflow, capable of reacting to external signals and executing durable, fault-tolerant operations, boosting your confidence in building production-ready systems.

Understanding Trigger.dev's Core: Events, Tasks, and Retries

Imagine you're building a system where different parts need to communicate and perform actions. A user signs up, a file is uploaded, or an external API sends a notification. How do you reliably kick off a series of operations in response to these occurrences, especially when those operations might take a long time or encounter temporary network issues? This is precisely where Trigger.dev's event-driven, durable execution model shines. It provides the primitives to build systems that are not just automated, but also fault-tolerant and scalable.

Events: The Spark of Your Workflow

At the heart of every Trigger.dev workflow is an **Event**. Think of an event as a notification or a signal that something significant has just happened in your system or an external service. It's the "spark" that kicks off your automated process.

- **What is an Event?** An event is a simple, structured JSON payload describing an occurrence. It's a lightweight, immutable record of something that happened at a specific point in time, like `user.signed.up` or `order.shipped`.

- **Why do Events matter?** Events enable highly decoupled systems. Instead of tightly coupling components by directly calling functions, parts of your system can simply emit an event. Any workflow interested in that event can then react to it, without needing to know who or what generated it. This decoupling is crucial for building scalable, flexible, and easily maintainable architectures, especially as your system grows and integrates with more services.
    - `⚡ Real-world insight:` In production, events are often used to bridge microservices or integrate with third-party webhooks (e.g., Stripe payments, GitHub pushes).
- **How do Events work in Trigger.dev?** You define a specific event name (e.g., `user.signed.up`, `invoice.processed`). When this event is sent to Trigger.dev (either via its SDK, API, or dashboard), any workflow configured to listen for that event will be triggered. Trigger.dev acts as the central nervous system, routing these events to the correct workflows.

Tasks: The Durable Workhorses

Once an event triggers a workflow, the actual work is performed by **Tasks**. A task is a single, self-contained unit of work that Trigger.dev executes reliably.

- **What is a Task?** In Trigger.dev, a task is essentially an asynchronous function that you define within your workflow using `io.runTask()`. It represents a specific operation, like sending an email, calling an external API, performing a complex calculation, or interacting with an AI model.
- **Why do Tasks matter?** Tasks are the core of Trigger.dev's durable execution model. This means that once a task starts, Trigger.dev ensures it either completes successfully or, in case of failure, retries it until it succeeds (or exhausts its retry attempts). This is crucial for long-running operations or those that interact with external, potentially unreliable services.
    - `🧠 Important:` Your code doesn't need to worry about saving its state or re-running from scratch if the server crashes or the network blips. Trigger.dev handles that for you by checkpointing the workflow's progress. This means your workflow can literally pause, be moved to another server, and resume exactly where it left off, making it incredibly robust.

- **How do Tasks work?** You define a task using `await io.runTask('your-task-id', async (payload, step) => { ... });`. The `io` object provides the context for running tasks and interacting with Trigger.dev's features, while the `step` context provides powerful capabilities for advanced workflows, which we'll explore later. For now, think of `io.runTask` as the reliable wrapper for your business logic.

Retries: Embracing Failure Gracefully

In distributed systems, failures are not exceptions; they are an expected part of daily operations. Network glitches, API rate limits, or temporary service outages can all cause an operation to fail. Instead of crashing, your workflows need to be resilient and self-healing. This is where **Retries** come in.

- **What are Retries?** Retries are the automatic re-execution of a failed task. Trigger.dev automatically retries tasks that encounter transient errors, giving them another chance to succeed without any manual intervention.
- **Why do Retries matter?** Retries are fundamental to building robust systems. They allow your workflows to recover from temporary issues, significantly increasing reliability and reducing operational burden. Without retries, a single momentary network blip could halt an entire critical workflow, requiring costly manual restarts.
  - `⚠️ What can go wrong:` While powerful, retries must be used thoughtfully. If a task is not idempotent (meaning running it multiple times with the same input has the same effect as running it once), you need to design your task carefully to avoid unintended side effects during retries (e.g., sending the same email multiple times).
- **How do Retries work?** Trigger.dev has built-in, intelligent retry mechanisms. By default, it will retry tasks with an exponential backoff strategy, meaning it waits longer between each subsequent attempt, reducing load on the failing service. You can also configure the maximum number of attempts (`maxAttempts`) and the retry delay (`retryDelayInMs`) for specific tasks to fine-tune this behavior.

Let's visualize how these core components interact to form a resilient workflow:

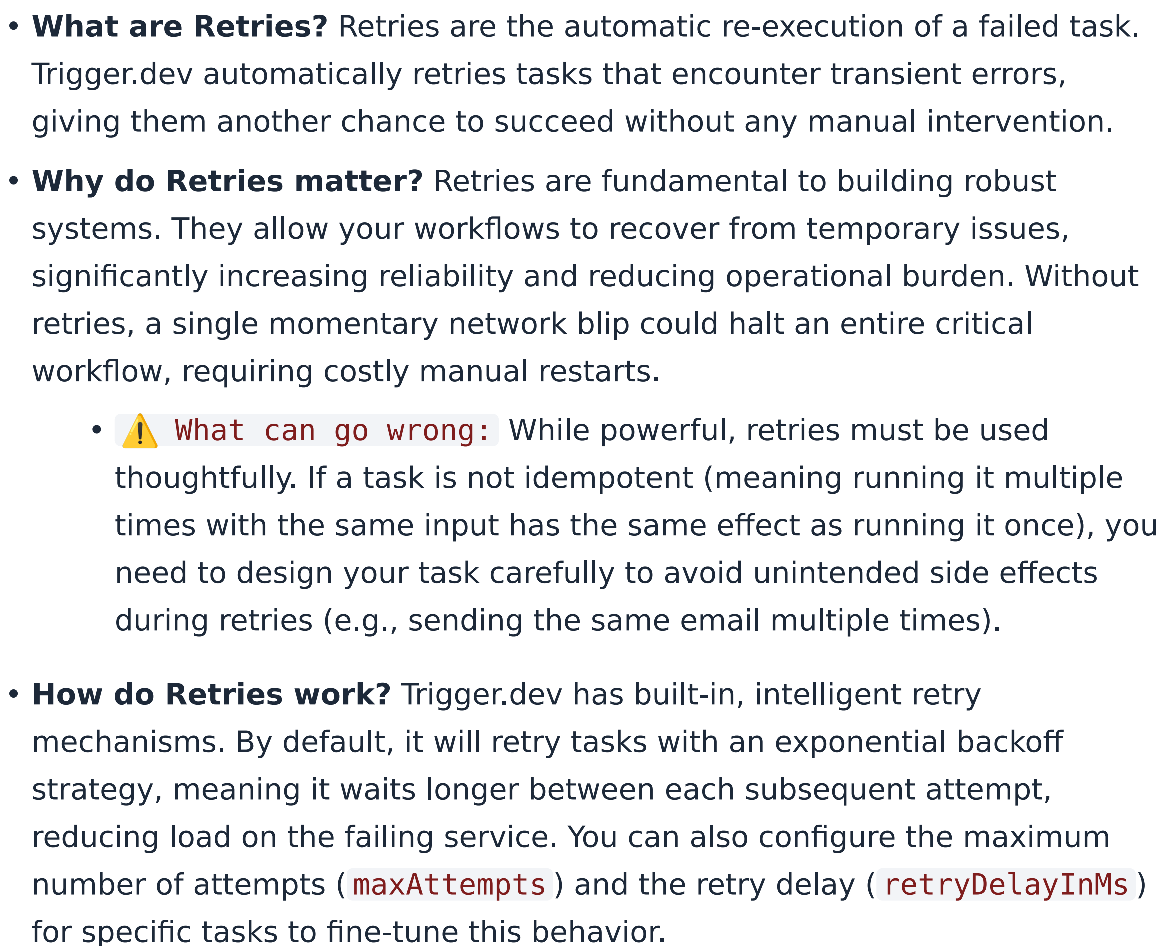

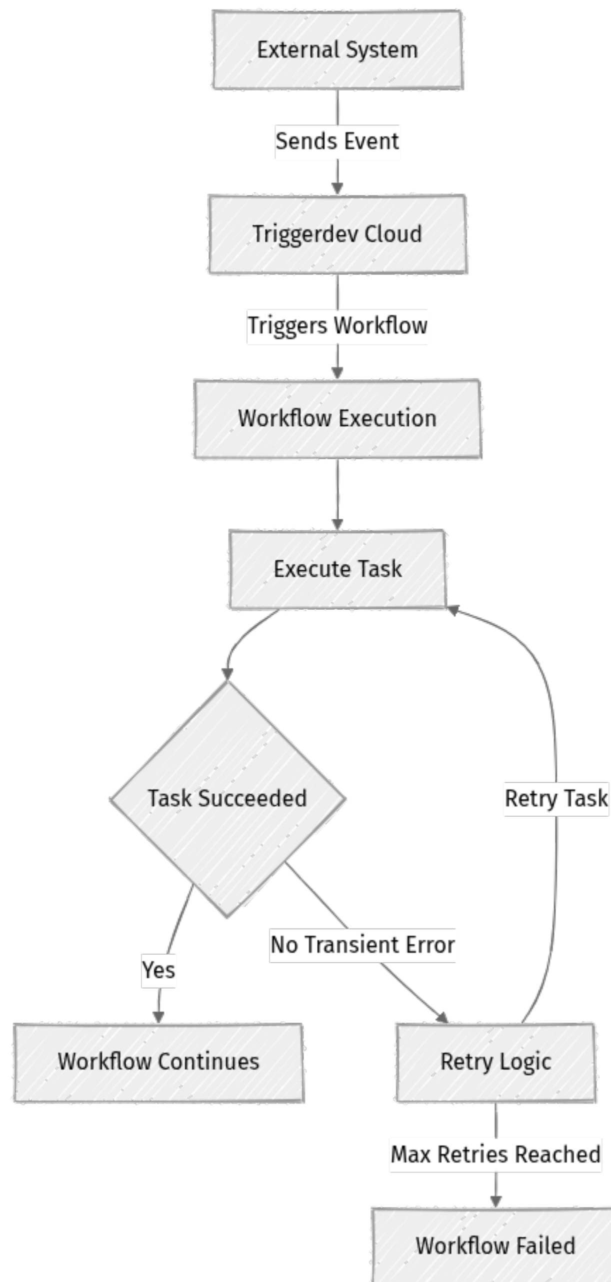

This diagram illustrates the flow: an `External System` sends an event, which `Trigger.dev Cloud` receives. This `Triggers Workflow Execution`, leading to `Execute Task`. If the `Task Succeeded?` is no due to a transient error, `Retry Logic` kicks in, potentially retrying the task with backoff. If `Max Retries Reached`, the `Workflow Failed`, allowing for alerts. If the task succeeds, the `Next Task or Workflow Continues`.

Your First Resilient Workflow: A Step-by-Step Build

Let's put these concepts into practice. We'll create a simple workflow that listens for a custom event, performs a "processing" task, and intentionally demonstrates how retries work to recover from a temporary failure.

First, ensure your Trigger.dev development server is running from Chapter 2. If not, navigate to your project directory and run:

```
npm run dev
```

This command starts your local Trigger.dev development server and the background process that watches for changes in your `src/jobs` directory.

Step 1: Create Your Workflow File

Inside your `src/jobs` directory, create a new file named `my-first-workflow.ts`. This is where we'll define our event listener and tasks.

Now, let's add the basic structure of a Trigger.dev job to `src/jobs/my-first-workflow.ts`:

```
// src/jobs/my-first-workflow.ts
import { client } from "../trigger";

client.defineJob({
  id: "my-first-workflow",
  name: "My First Workflow",
  version: "1.0.0",
  // We'll define the trigger and run function here shortly
});
```

- `import { client } from "../trigger";`: This line imports the pre-configured Trigger.dev client instance that we set up in Chapter 2. This client is your gateway to defining jobs, sending events, and interacting with the Trigger.dev platform.
- `client.defineJob({...});`: This is the main function you use to define a workflow (or "job" in Trigger.dev terminology).
  - `id`: A unique, machine-readable identifier for your workflow. Good practice is to use kebab-case.
  - `name`: A human-readable name that will appear in the Trigger.dev dashboard.
  - `version`: A semantic version for your workflow, useful for tracking changes over time.

Next, let's add the `trigger` and `run` properties to our job definition. The `trigger` specifies when the workflow should run, and `run` specifies what the workflow should do.

```
// src/jobs/my-first-workflow.ts
import { client } from "../trigger";

client.defineJob({
  id: "my-first-workflow",
  name: "My First Workflow",
  version: "1.0.0",
  // This is where we define what event triggers this workflow
  trigger: client.on("my.event", {
    schema: {
      type: "object",
      properties: {
        message: { type: "string" },
        shouldFail: { type: "boolean" },
      },
      required: ["message", "shouldFail"],
      additionalProperties: false,
    },
  }),
  // This is the core function that runs when the event is received
  run: async (payload, io, ctx) => {
    // We'll add our tasks here in the next step
  },
});
```

Let's break down these additions:

- `trigger: client.on("my.event", { schema: ... })`: This tells Trigger.dev to listen for an event named `my.event`.
    - `client.on("my.event", ...)`: This is how you subscribe your workflow to a specific event.
    - `schema`: This object defines the expected structure of the incoming event payload using JSON Schema. This is excellent for early validation and type safety. Here, we're expecting a `message` (string) and `shouldFail` (boolean). If an incoming event doesn't match this schema, Trigger.dev will reject it, preventing unexpected errors in your workflow.

- `run: async (payload, io, ctx) => { ... };`: This is the heart of your workflow, an asynchronous function that executes when `my.event` is received.
  - `payload`: This argument contains the data from the `my.event` that triggered this workflow. Its type is automatically inferred from your `schema`.
  - `io`: This is the I/O client, providing methods for performing operations like logging (`io.logger`), running durable tasks (`io.runTask`), and interacting with other Trigger.dev features.
  - `ctx`: This is the context object, providing information about the current execution, including details about retries (`ctx.attempt`), environment, and more.

Finally, let's add the actual tasks inside the `run` function, demonstrating the core concepts of tasks and retries.

```
// src/jobs/my-first-workflow.ts
import { client } from "../trigger";

client.defineJob({
  id: "my-first-workflow",
  name: "My First Workflow",
  version: "1.0.0",
  trigger: client.on("my.event", {
    schema: {
      type: "object",
      properties: {
        message: { type: "string" },
        shouldFail: { type: "boolean" },
      },
      required: ["message", "shouldFail"],
      additionalProperties: false,
    },
  }),
  run: async (payload, io, ctx) => {
    // Log the incoming event payload for observability
    io.logger.info("Received 'my.event' payload", payload);

    // Step 1: Simulate a data processing task that might fail temporarily
    await io.runTask(
      "process-incoming-data", // Unique ID for this specific task within the
workflow
      async (taskPayload) => {
        io.logger.info(`Processing data: ${taskPayload.message}`);

        // Introduce an intentional, temporary failure for demonstration
        // This task will fail if `shouldFail` is true AND it's the first
attempt (attempt.number is 1).
        // It will succeed on the second attempt, demonstrating retry
recovery.
        if (taskPayload.shouldFail && ctx.attempt.number === 1) {
          io.logger.warn(`Simulating a temporary failure for attempt ${ctx.att
```

```
empt.number}...`);
          throw new Error("Simulated transient error during processing!");
        }

        io.logger.info("Data processed successfully!");
        return { status: "processed", originalMessage: taskPayload.message };
      },
      // The payload for this specific task (can be different from the event
payload)
      payload,
      // Configure retries for this task
      {
        maxAttempts: 3, // Try up to 3 times
        retryDelayInMs: 1000, // Wait 1 second between retries
      }
    );

    // Step 2: Simulate another task that depends on the first one
    // This task will only run if the previous 'process-incoming-data' task
succeeds (potentially after retries).
    await io.runTask("send-confirmation", async (taskPayload) => {
      io.logger.info(`Sending confirmation for: ${taskPayload.originalMessage}
`);
      // In a real application, you might send an email, update a database,
notify another service, etc.
      return { confirmationSent: true };
    });

    io.logger.info("Workflow completed successfully!");
  },
});
```

Here's a detailed explanation of the new code within the `run` function:

- `io.logger.info(...)`: This is Trigger.dev's structured logger. Using `io.logger` ensures your logs are captured by Trigger.dev and visible in the dashboard, making debugging much easier than relying solely on `console.log`.

- `await io.runTask(...)`: This is how you define and execute a durable task.
    - `"process-incoming-data"`: This is a unique identifier for this specific task within your workflow. This ID is crucial for Trigger.dev to track the task's state, enable retries, and ensure durable execution.
    - `async (taskPayload) => { ... }`: This is the function containing the actual logic for your task. It receives a `taskPayload` (which we're passing the main event `payload` to) and can perform any asynchronous operations.
    - `payload`: This is the data we're passing to our `process-incoming-data` task. In this case, we're simply forwarding the entire `payload` that triggered the workflow.
    - `{ maxAttempts: 3, retryDelayInMs: 1000 }`: This is the retry configuration for this specific task.
        - `maxAttempts: 3`: Tells Trigger.dev to try this task up to 3 times if it fails.
        - `retryDelayInMs: 1000`: Specifies a 1-second delay between retry attempts. Trigger.dev often applies exponential backoff by default, even if you specify a fixed delay, to prevent overwhelming a failing service.
    - `if (taskPayload.shouldFail && ctx.attempt.number === 1)`: This line is a clever way to simulate a transient failure.
        - `taskPayload.shouldFail`: This condition comes directly from the event payload.
        - `ctx.attempt.number === 1`: The `ctx.attempt.number` tells you which attempt the current task execution is. It starts at `1` for the first attempt. By checking `=== 1`, we ensure the task only fails on its very first execution. On subsequent retries (attempt 2, 3, etc.), this condition will be false, and the task will succeed, demonstrating how retries gracefully handle temporary issues.
- `await io.runTask("send-confirmation", ...)`: This defines a second task. Notice that this task is placed after the first `io.runTask`. Due to Trigger.dev's durable execution, this `send-confirmation` task will only start once `process-incoming-data` has successfully completed (which might involve one or more retries). This sequential execution is guaranteed.

Step 2: Trigger Your Workflow

With your Trigger.dev development server running (via `npm run dev`), it automatically detects changes to `src/jobs` and registers your new workflow. You can now trigger it!

There are a few ways to send an event to Trigger.dev, but for local development, the dashboard is often the easiest.

1. **Using the Trigger.dev Dashboard (Recommended for local dev):**
    - Open your browser to the Trigger.dev dashboard, usually at `<http://localhost:8080>`.
    - Navigate to the "Events" section on the left sidebar.
    - Click the "Send Event" button (usually in the top right).
    - For the "Event Name", type `my.event` (this must exactly match the `client.on("my.event", ...)` you defined).
    - For the "Payload (JSON)", enter the following JSON. This payload includes `shouldFail: true` to demonstrate the retry mechanism.

```
        {
          "message": "Hello from Trigger.dev!",
          "shouldFail": true
        }
```

```
- Click "Send Event".
```

1. **Using `curl` (for quick API testing):** Open a new terminal window and run the following command. Remember to replace `<YOUR_DEV_API_KEY>` with the `TRIGGER_API_KEY` found in your `.env` file from Chapter 2.

```
    curl -X POST http://localhost:8080/api/v1/events \
         -H "Content-Type: application/json" \
         -H "Authorization: Bearer <YOUR_DEV_API_KEY>" \
         -d '{ "name": "my.event", "payload": { "message": "Hello via curl!",
"shouldFail": true } }'
```

Step 3: Observe Execution and Retries

After sending the event, switch back to your Trigger.dev dashboard (`<http://localhost:8080>`).

- Go to the "Runs" section on the left sidebar. You should see a new run entry for "My First Workflow".

- Click on the run ID to see its detailed execution timeline.
- You'll observe the `process-incoming-data` task initially failing (its status will briefly show "Retrying"), then pausing for 1 second, and finally succeeding on the second attempt.
- After `process-incoming-data` successfully completes, the `send-confirmation` task will then execute successfully.
- Check your development server's console output (where `npm run dev` is running). You'll see the `io.logger.info` messages and critically, the `io.logger.warn` message during the simulated failure, followed by the successful log on the retry.

This hands-on experience clearly demonstrates:

- How an `event` (`my.event`) kicks off a workflow.
- How `tasks` (`process-incoming-data`, `send-confirmation`) encapsulate durable units of work.
- How Trigger.dev automatically `retries` failed tasks, making your workflow resilient to transient errors.

Mini-Challenge: Enhancing Your Workflow

Now it's your turn to make a small modification! This challenge will reinforce your understanding of event payloads and conditional logic within tasks.

Challenge: Modify the `my-first-workflow.ts` file.

1. **Add a new property** to the `my.event` schema called `priority` (type `string`, can be "high" or "low"). Make it a required property.
2. **Update the `run` function** to check the `priority` from the incoming `payload`.
3. If `priority` is "high", make the `process-incoming-data` task simulate a failure twice (meaning it should succeed on the third attempt).
4. If `priority` is "low", the `process-incoming-data` task should never fail, regardless of the `shouldFail` flag.

Hint:

- Remember to update the `schema` in `client.on()` to include `priority`.
- Inside `io.runTask`, you can use `taskPayload.priority` and `ctx.attempt.number` to control the simulated failure logic.

- Test with new payloads from the dashboard:
    - High priority, fails twice: `{"message": "High priority data", "shouldFail": true, "priority": "high"}`
    - Low priority, never fails: `{"message": "Low priority data", "shouldFail": true, "priority": "low"}`
    - Low priority, never fails (even if `shouldFail` is true): `{"message": "Another low priority", "shouldFail": true, "priority": "low"}`

What to observe/learn: This challenge reinforces how to use event payload data to dynamically alter workflow behavior and further demonstrates the power of `ctx.attempt.number` in managing retry logic for different scenarios. You'll also practice modifying your event schema and observing its impact on workflow execution.

Common Pitfalls & Troubleshooting Basic Workflows

Even with simple workflows, a few common issues can arise. Knowing how to spot and fix them will save you significant time and frustration.

- **Mismatched Event Payloads:**
    - **Pitfall:** You send an event payload that doesn't match the `schema` defined in `client.on()`. This can lead to validation errors, prevent your workflow from triggering, or result in `undefined` values in your `payload` within the `run` function.
    - **Troubleshooting:**

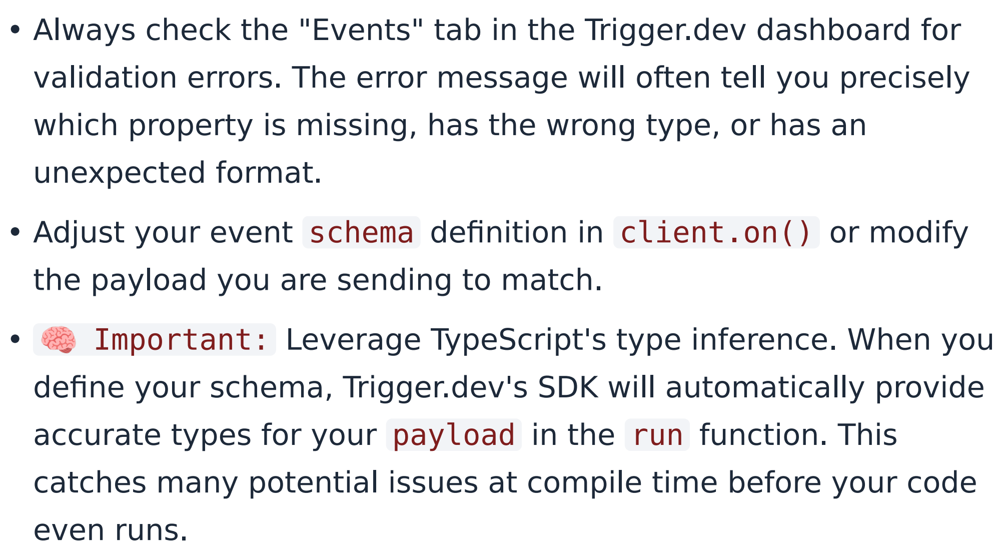

        - Always check the "Events" tab in the Trigger.dev dashboard for validation errors. The error message will often tell you precisely which property is missing, has the wrong type, or has an unexpected format.
        - Adjust your event `schema` definition in `client.on()` or modify the payload you are sending to match.
        - `🧠 Important:` Leverage TypeScript's type inference. When you define your schema, Trigger.dev's SDK will automatically provide accurate types for your `payload` in the `run` function. This catches many potential issues at compile time before your code even runs.

- **Infinite Retries (or Too Few):**
    - **Pitfall:** A task might get stuck in an infinite retry loop if `maxAttempts` is set too high for a persistent error (e.g., an incorrect API key that will never work), or conversely, a task might fail permanently too quickly if `maxAttempts` is too low for a truly transient error.
    - **Troubleshooting:**

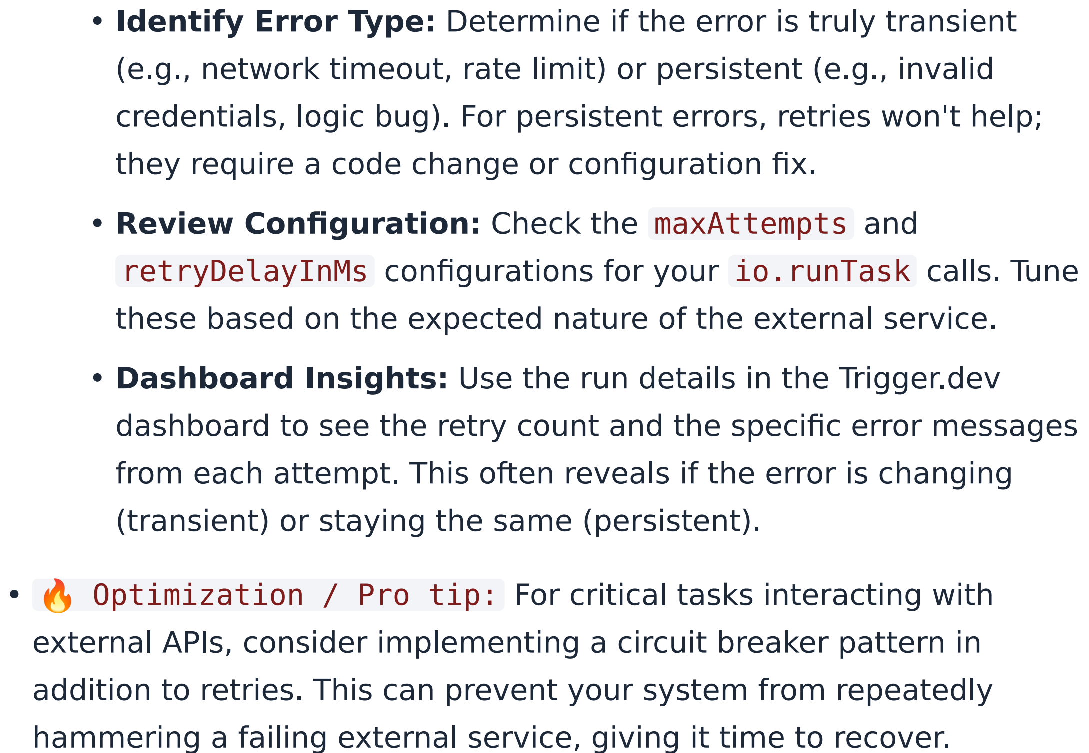

        - **Identify Error Type:** Determine if the error is truly transient (e.g., network timeout, rate limit) or persistent (e.g., invalid credentials, logic bug). For persistent errors, retries won't help; they require a code change or configuration fix.
        - **Review Configuration:** Check the `maxAttempts` and `retryDelayInMs` configurations for your `io.runTask` calls. Tune these based on the expected nature of the external service.
        - **Dashboard Insights:** Use the run details in the Trigger.dev dashboard to see the retry count and the specific error messages from each attempt. This often reveals if the error is changing (transient) or staying the same (persistent).
    - `🔥 Optimization / Pro tip:` For critical tasks interacting with external APIs, consider implementing a circuit breaker pattern in addition to retries. This can prevent your system from repeatedly hammering a failing external service, giving it time to recover.

- **Local Development Quirks:**
    - **Pitfall:** Your workflow isn't being triggered, or recent code changes aren't reflected in your running application.
    - **Troubleshooting:**
        - **`npm run dev` Status:** Ensure your `npm run dev` process is still running in your terminal. If it crashed, restart it.
        - **Terminal Logs:** Check the terminal where `npm run dev` is running for any errors or messages about job registration. Trigger.dev should log when it successfully registers your `my-first-workflow.ts` file.
        - **Environment Variables:** Verify that your `TRIGGER_API_KEY` and `TRIGGER_PUBLIC_KEY` in your `.env` file are correct and haven't been accidentally changed. These are essential for your local client to connect to the Trigger.dev dev server.
        - **Restart:** Sometimes, simply stopping (`Ctrl+C`) and restarting `npm run dev` can resolve issues with file change detection or stale configurations.

Summary: Your Workflow Foundation

Congratulations! You've just built and observed your first resilient workflow with Trigger.dev. We covered three crucial concepts that form the very foundation of event-driven, durable systems:

- **Events** act as the triggers, initiating your workflows based on external or internal signals, enabling loose coupling.
- **Tasks** are the durable units of work, executed reliably by Trigger.dev, ensuring your operations complete even through transient failures by preserving state and resuming execution.
- **Retries** provide built-in fault tolerance, automatically re-attempting failed tasks with intelligent backoff to recover from temporary issues, significantly improving system reliability.

Understanding these fundamentals is key to building any automation or AI agent with Trigger.dev. You now have a solid foundation for creating event-driven, robust applications that can withstand the challenges of production environments.

In the next chapter, we'll explore more advanced workflow patterns, including how to schedule tasks for future execution and manage dependencies between different steps. Get ready to add another layer of sophistication to your Trigger.dev skills!

References

- Trigger.dev Documentation: Events
- Trigger.dev Documentation: Tasks
- Trigger.dev Documentation: Retries
- Trigger.dev GitHub Repository

This page is AI-assisted and reviewed. It references official documentation and recognized resources where relevant.

CHAPTER 04

Building Robust Workflows: Queues, Scheduling, and Long-Running Processes

In the world of modern applications, especially those involving AI agents or complex data processing, tasks often need to run reliably in the background, at specific times, or endure for extended periods without interruption. Imagine sending out millions of personalized emails, generating daily reports, or orchestrating a multi-step AI inference process. How do you ensure these operations complete successfully, even if your server crashes or an external API temporarily fails?

This chapter dives deep into the core mechanisms Trigger.dev provides to build such resilient systems: **queues**, **scheduling**, and **long-running durable execution**. We'll learn how these concepts work together to create workflows that are not just functional, but truly robust and production-ready. By the end, you'll be equipped to design and implement background tasks that can handle failures gracefully, execute on a precise timetable, and manage complex, multi-stage operations with ease.

If you've followed the previous chapters, you've already seen how to define basic jobs. Now, we'll enhance those jobs with advanced capabilities that are crucial for any serious application.

The Pillars of Robustness: Queues, Scheduling, and Durable Execution

Building systems that reliably perform tasks, especially when those tasks are asynchronous, time-sensitive, or long-lived, requires a solid foundation. Trigger.dev provides this foundation by abstracting away much of the complexity associated with distributed systems. It allows developers to focus on the business logic rather than the intricate details of fault tolerance and state management.

What are Queues and Why Do We Need Them?

Imagine a popular e-commerce site on Black Friday. Thousands of orders are coming in every second. If each order immediately tried to process payment, update inventory, and send confirmation emails synchronously, the system would quickly buckle under the load. This is where queues come in.

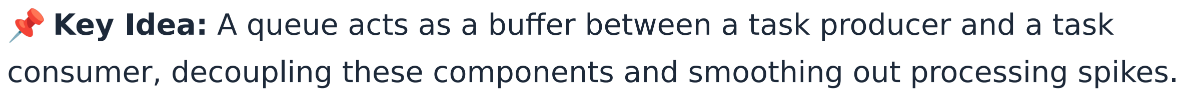

📌 **Key Idea:** A queue acts as a buffer between a task producer and a task consumer, decoupling these components and smoothing out processing spikes.

When you send an event to Trigger.dev, it doesn't immediately execute your job. Instead, the event (representing a task) is placed into a managed queue. Your Trigger.dev worker then picks up events from this queue and processes them at a controlled pace.

Why Queues Matter for Production Systems:

- **Decoupling:** The part of your application that creates a task (e.g., a user clicking "checkout") doesn't need to know how or when the task will be completed. It just sends it to the queue. This improves system architecture by reducing direct dependencies.
- **Load Leveling:** Prevents your backend from being overwhelmed during traffic spikes. Tasks are processed at a manageable rate, preventing system crashes and ensuring consistent performance.
- **Reliability:** If a worker fails while processing a task, the task can often be returned to the queue and retried by another worker, ensuring it eventually completes. This is critical for maintaining data integrity and business continuity.
- **Scalability:** You can easily scale your workers up or down independently of the rate at which tasks are produced. This allows for efficient resource allocation based on current demand.

In Trigger.dev, when you `client.sendEvent()`, that event essentially enters a managed queue, ready for your defined job to pick it up. This abstraction simplifies building asynchronous systems significantly.

Scheduling Tasks: Doing Things on Time

Some tasks aren't triggered by an immediate event but need to happen at specific intervals or times. Think of:

- Generating a daily sales report at 9:00 AM.
- Sending out weekly newsletters every Monday morning.
- Checking for expired user sessions every hour.
- Running a nightly database backup or data synchronization.

Trigger.dev supports scheduling using standard **Cron expressions**. If you've ever set up a cron job on a Linux server, you'll be familiar with the syntax. It allows you to define highly precise, recurring schedules for your jobs. This eliminates the need for external cron services or complex server-side task managers.

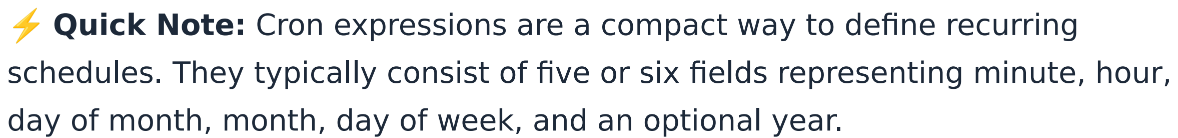

⚡ **Quick Note:** Cron expressions are a compact way to define recurring schedules. They typically consist of five or six fields representing minute, hour, day of month, month, day of week, and an optional year.

Durable Execution for Long-Running Workflows

Consider an AI agent workflow that involves multiple, potentially time-consuming steps:

1. Transcribing a long audio file (5 minutes).
2. Summarizing the transcription using an LLM (2 minutes).
3. Generating an image based on the summary (1 minute).
4. Sending the result for human review (human might take 30 minutes to 2 hours).
5. Publishing the final output (1 minute).

This entire process could easily take an hour or more. What happens if your server restarts in the middle of transcription, or the LLM API times out? Without durable execution, the entire workflow might fail, losing all progress and requiring a complete restart, which is inefficient and costly.

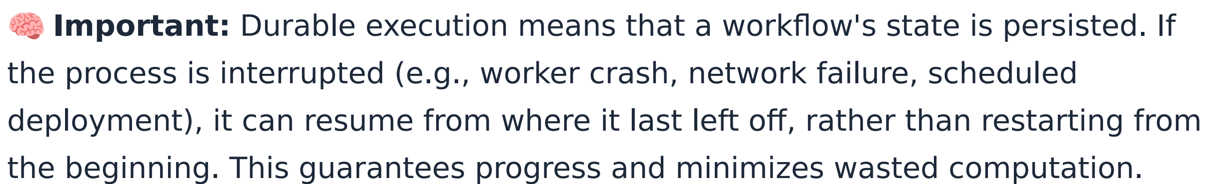

🧠 **Important:** Durable execution means that a workflow's state is persisted. If the process is interrupted (e.g., worker crash, network failure, scheduled deployment), it can resume from where it last left off, rather than restarting from the beginning. This guarantees progress and minimizes wasted computation.

Trigger.dev achieves durable execution by checkpointing the state of your job. When your job encounters an `await` or needs to pause (e.g., for a delay, a retry, or waiting for an external event), Trigger.dev saves its current state. If the worker process dies, another available worker can pick up the job instance and resume it from the last saved state. This is incredibly powerful for:

- **Long-running tasks:** No more worrying about server restarts, temporary network issues, or deployment cycles interrupting critical operations.
- **Human-in-the-loop workflows:** Your workflow can pause indefinitely, waiting for human input or approval, and then resume seamlessly.
- **Retries:** If an API call fails, Trigger.dev can automatically retry it after a delay, ensuring forward progress without manual intervention.

Step-by-Step Implementation: Building Robust Jobs

Let's put these concepts into practice. We'll create a Trigger.dev project (or continue from previous chapters) and implement jobs that leverage queues, scheduling, and durable execution.

Prerequisites

Ensure you have your Trigger.dev project set up. If not, you can quickly initialize one using the latest `v4-beta` version:

```
npx trigger.dev@v4-beta init
```

Choose the Next.js option and follow the prompts. Make sure your environment variables (`TRIGGER_SECRET_KEY`, `TRIGGER_PUBLIC_KEY`) are correctly configured as per Chapter 2. This setup ensures your local development environment can connect to the Trigger.dev cloud service.

1. A Simple Queued Job Example

Every job you define in Trigger.dev is inherently queued. When you trigger an event, it goes into a queue. Let's define a job that simulates a background processing task, such as image manipulation.

Open your `src/trigger/jobs.ts` (or equivalent) file and add the following code. This file is where all your Trigger.dev jobs are defined.

```
// src/trigger/jobs.ts
import { client } from "./client";
import { eventTrigger } from "@trigger.dev/sdk";

// Define a job that simulates processing an image
client.defineJob({
  id: "process-image-job", // A unique identifier for this job
  name: "Process Image in Background", // A human-readable name for the
dashboard
  version: "1.0.0", // Helps manage job changes and deployments
  // This job is triggered by an event with a 'process.image' name
  trigger: eventTrigger({
    name: "process.image",
    // Define the expected structure of the event payload for type safety
    schema: {
      url: "string", // The URL of the image to process
      userId: "string", // The ID of the user who uploaded the image
    },
  }),
  // The 'run' function contains the core logic of your job
  run: async (payload, io, ctx) => {
    // Log the received payload for debugging and observability
    await io.logger.info("Starting image processing...", { payload });
```

```
    // Simulate a long-running image processing task.
    // The `io.wait` function is crucial for durable execution.
    // If the worker process restarts during these 3 seconds, the job will
    // automatically resume from this exact point on another available worker.
    await io.wait("3 seconds");

    // In a real scenario, you'd integrate with an actual image processing
service here.
    // For example:
    // const processedImageResult = await
someImageService.process(payload.url);
    // await io.logger.info("Image service responded", { result:
processedImageResult });

    await io.logger.info(`Image ${payload.url} processed for user ${payload.us
erId}.`);

    // You could send another event here to notify the user or trigger another
job.
    // For example:
    // await io.sendEvent("image.processed", { userId: payload.userId,
processedUrl: "..." });

    return { message: "Image processing complete!" };
  },
});
```

Let's break down the key elements of this job definition:

- `id`, `name`, `version`: These are standard metadata. The `id` is crucial as it uniquely identifies your job within Trigger.dev.
- `trigger: eventTrigger({ name: "process.image", schema: { ... } })`: This tells Trigger.dev that this job should run whenever an event named `"process.image"` is received. The `schema` defines the expected data structure for the `payload`, providing valuable type-checking and documentation.
- `run: async (payload, io, ctx) => { ... }`: This is the asynchronous function containing your job's logic.
    - `payload`: An object containing the data sent with the triggering event.
    - `io`: The Trigger.dev I/O client. This object provides durable operations like logging (`io.logger`), pausing (`io.wait`), and interacting with external services in a retryable manner (`io.runTask`). **All operations performed with `io` are durable.**
    - `ctx`: Provides context about the current job run, such as the run ID.

- `await io.wait("3 seconds")`: This is a fundamental building block of durable execution. Unlike `setTimeout`, which is non-durable and would lose state on a worker restart, `io.wait` tells Trigger.dev to pause the job, persist its state, and then resume it after the specified duration.

Triggering the Queued Job

You can trigger this job from your Next.js API route or anywhere you have access to the `client` instance. Let's create a new API route, for example, `src/app/api/trigger-image-process/route.ts`, that will send the `process.image` event.

```
// src/app/api/trigger-image-process/route.ts
import { client } from "@/trigger/client"; // Adjust path based on your
project structure
import { NextResponse } from "next/server";

// This Next.js API route will handle POST requests
export async function POST(request: Request) {
  const { imageUrl, userId } = await request.json();

  // Basic validation for incoming data
  if (!imageUrl || !userId) {
    return NextResponse.json({ error: "Missing imageUrl or userId" }, {
status: 400 });
  }

  // Send the event to Trigger.dev. This event will be queued, and the
  // 'process-image-job' we defined earlier will pick it up for execution.
  const event = await client.sendEvent({
    name: "process.image", // The name of the event this job listens for
    payload: {
      url: imageUrl,
      userId: userId,
    },
  });

  // Respond to the client indicating the event was successfully sent
  return NextResponse.json({
    message: "Image processing event sent!",
    eventId: event.id, // The ID of the event in Trigger.dev
  });
}
```

Now, if you send a `POST` request to `/api/trigger-image-process` (e.g., using `curl`, Postman, or a frontend fetch call) with a JSON body like `{"imageUrl": "https://example.com/pic.jpg", "userId": "user123"}`, Trigger.dev will receive the event, queue it, and your worker will eventually process it. You can observe the job's status and logs in the Trigger.dev dashboard, seeing it pause for 3 seconds and then complete.

2. Implementing a Scheduled Job

Next, let's create a job that runs on a predefined schedule, rather than in response to an event. This is perfect for periodic tasks like health checks, data synchronization, or report generation.

Add this job definition to your `src/trigger/jobs.ts` file, alongside your `process-image-job`:

```
// src/trigger/jobs.ts (add to existing file)
import { client } from "./client";
import { cronTrigger } from "@trigger.dev/sdk"; // Import cronTrigger

// Define a job that runs on a schedule
client.defineJob({
  id: "scheduled-health-check", // Unique ID for this scheduled job
  name: "Hourly System Health Check",
  version: "1.0.0",
  // This job is triggered by a cron schedule
  trigger: cronTrigger({
    // Cron expression for every minute: "minute hour dayOfMonth month
dayOfWeek"
    // For demonstration, let's run it every minute: "* * * * *"
    // For every hour at minute 0: "0 * * * *"
    cron: "* * * * *",
  }),
  run: async (payload, io, ctx) => {
    await io.logger.info("Running system health check...");

    // Simulate checking various system components
    const status = {
      database: "healthy",
      api_gateway: "healthy",
      cache: "unhealthy", // Oh no, a problem!
    };

    if (status.cache === "unhealthy") {
      await io.logger.error("Cache system is unhealthy!", { status });
      // In a real scenario, you might send an alert or trigger a remediation
job.
      // await io.sendEvent("alert.system.cache.unhealthy", { status });
    } else {
      await io.logger.info("All systems are nominal.", { status });
    }

    return { message: "Health check complete", status };
  },
});
```

- `cronTrigger({ cron: "* * * * *" })`: This is how you define a scheduled job. The `cron` property takes a standard cron expression. `* * * * *` means "every minute of every hour of every day of every month of every day of the week."

Cron Expression Basics:

A cron expression typically consists of 5 fields, representing the schedule:```

| | | | | | | | | ----- Day of week (0 - 7, Sunday is 0 or 7) | | | ------- Month (1 - 12) | | --------- Day of month (1 - 31) | ----------- Hour (0 - 23) ------------- Minute (0 - 59)

```
**Common Examples:**
- `0 * * * *`: Every hour at the 0th minute (e.g., 1:00, 2:00, etc.)
- `0 9 * * 1`: Every Monday at 9:00 AM
- `0 0 1 * *`: On the first day of every month at midnight (midnight on the
1st)
- `*/5 * * * *`: Every 5 minutes

Once you deploy your Trigger.dev worker (by running your Next.js app in
production mode or deploying it), this job will automatically start running
according to its schedule. You'll see new job runs appearing in your
Trigger.dev dashboard every minute (or as per your cron expression).

### 3. Combining Durability with Retries for Reliability

Trigger.dev jobs come with built-in retry mechanisms, which are essential for
handling transient failures (e.g., a temporary network glitch or an external
API being briefly unavailable). When an `io` operation (like `io.runTask` or
an external API call wrapped in `io.runTask`) throws an error, Trigger.dev can
automatically retry the step. This is a huge advantage over traditional
background task systems where you'd have to implement complex retry logic
yourself.

Let's modify our image processing job to include a simulated external API call
that might fail and show how retries work.
```typescript
// src/trigger/jobs.ts (modify the existing 'process-image-job')
import { client } from "./client";
import { eventTrigger } from "@trigger.dev/sdk";

client.defineJob({
  id: "process-image-job",
  name: "Process Image in Background",
  version: "1.0.1", // Increment version since we're changing the logic
  trigger: eventTrigger({
    name: "process.image",
    schema: {
      url: "string",
      userId: "string",
    },
  }),
  run: async (payload, io, ctx) => {
    await io.logger.info("Starting image processing...", { payload });

    // Simulate an external API call that might fail.
    // `io.runTask` wraps this logic, making it durable and retryable.
    const processedData = await io.runTask(
      "call-image-api", // Unique ID for this specific task step within the
job
      async () => {
        // Simulate a random failure for demonstration purposes.
        // In a real app, this would be an actual external API call.
        const shouldFail = Math.random() < 0.5; // 50% chance of failure
```
```

```
        if (shouldFail) {
          // If an error is thrown here, Trigger.dev will automatically retry
this step.
          throw new Error("Simulated external image API failure!");
        }

        // Simulate success and return some processed data.
        await io.logger.info("Successfully called external image API.");
        return {
          originalUrl: payload.url,
          processedUrl: `https://processed.example.com/${payload.userId}-$
{Date.now()}.jpg`,
          metadata: { width: 800, height: 600 },
        };
      },
      {
        // Optional: Configure specific retry options for this task step.
        // Trigger.dev automatically retries `io.runTask` calls that throw
errors
        // using a default exponential backoff strategy.
        // For example, to retry 3 times with a 10-second initial delay:
        // retries: { maxAttempts: 3, factor: 1, minTimeoutInMs: 10000 }
      }
    );

    // This part of the code will only execute if 'call-image-api' succeeds
    // (potentially after several retries).
    await io.logger.info(`Image processing complete for ${payload.url}.`, {
      processedData,
    });

    return { message: "Image processing complete!", data: processedData };
  },
});

// ... (keep the scheduled-health-check job below this if you defined it)
```

In this updated job:

- We incremented the `version` to `1.0.1`. It's a good practice to update the version whenever you make significant logic changes to a job, especially in production environments, to ensure proper deployment and potential rollback capabilities.
- `io.runTask("call-image-api", async () => { ... })`: This wraps our potentially failing logic. If the `async` function passed to `io.runTask` throws an error, Trigger.dev will automatically retry this specific step of the job, based on its default retry policy (usually exponential backoff). This means your entire workflow doesn't restart from the beginning; only the failing part is re-attempted.

- The retry options are commented out, showing where you could customize them. By default, Trigger.dev provides robust retry behavior, often sufficient for most transient failures.

When you trigger this job (via the API route), you might see it fail a few times and then retry, eventually succeeding, or failing permanently after exhausting its retry attempts. This behavior is fully observable in the Trigger.dev dashboard, providing a clear audit trail of each attempt.

Visualizing a Robust Workflow

This diagram illustrates how an event enters a queue, is picked up by a worker, and how `io.runTask` provides a durable, retryable step within the workflow, ensuring resilience against failures.

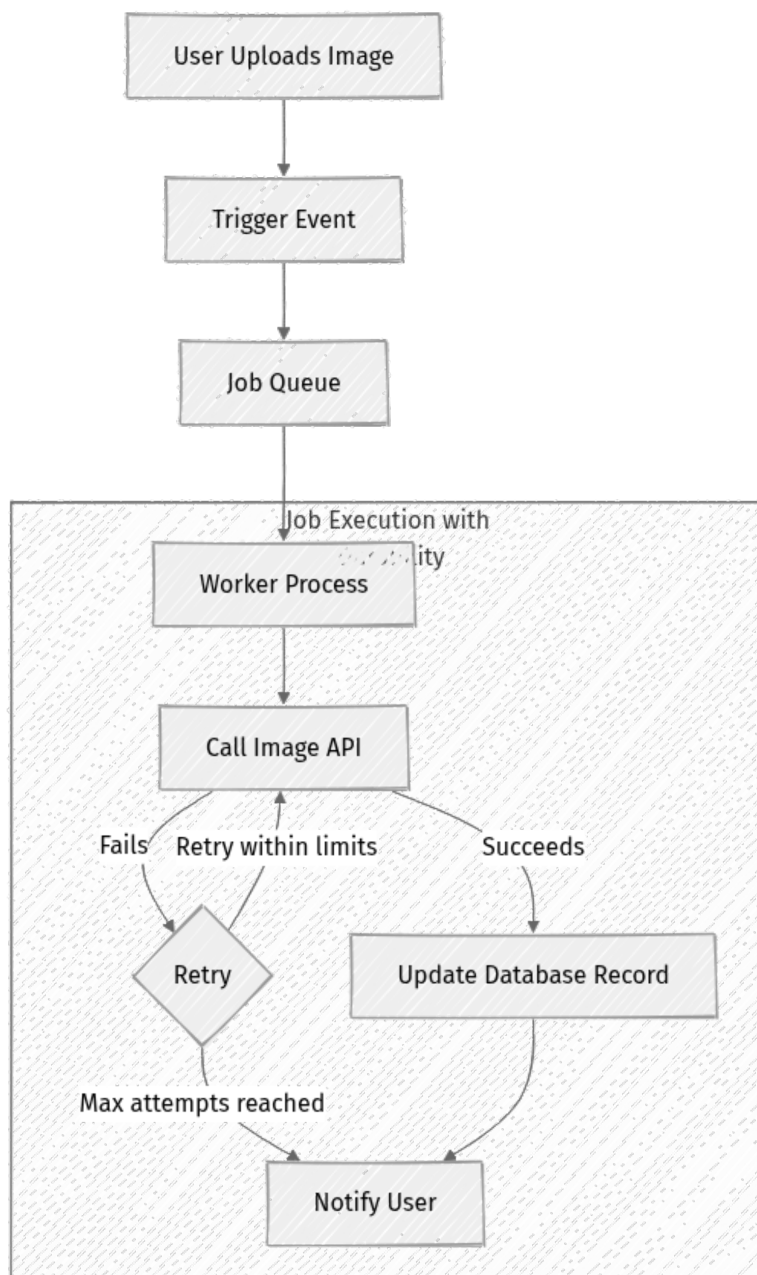

Mini-Challenge: Scheduled Report Generation

Now it's your turn to apply what you've learned! Create a new job that simulates generating an hourly report.

Challenge:

1. **Define a new Trigger.dev job** in your `src/trigger/jobs.ts` file.
2. **Schedule it to run once every hour**, specifically at 30 minutes past the hour (e.g., 1:30, 2:30, etc.).
3. **Inside the job's `run` function:**
    - Simulate fetching data from an external analytics API. You can use `io.wait("5 seconds")` to represent the API call's duration and introduce a `Math.random() < 0.6` (60% chance of failure) to simulate an unreliable API.
    - **Crucially, ensure this simulated API call is retryable.** If it fails, Trigger.dev should automatically retry it at least twice before finally giving up.
    - If the API call is successful, log a message using `io.logger.info` indicating a report was "generated" for the current hour, including the current timestamp (e.g., `new Date().toISOString()`).
    - If it ultimately fails after all retries, log an `io.logger.error` message.

Hint:

- For the cron expression, `30 * * * *` will run at 30 minutes past every hour.
- Use `io.runTask` for the simulated API call. You can explicitly set `retries: { maxAttempts: 3 }` (meaning 1 initial attempt + 2 retries) within the `io.runTask` options.
- Remember to give your job a unique `id` and `name`.

What to observe/learn:

- How scheduled jobs automatically appear in your Trigger.dev dashboard and execute at the precise times.
- How Trigger.dev handles retries for `io.runTask` calls, even across potential worker restarts, ensuring durability.
- The difference in dashboard logs and status for a job that completes successfully, fails after exhausting retries, or succeeds on a retry attempt.

Common Pitfalls & Troubleshooting

Even with robust tools like Trigger.dev, distributed systems can present unique challenges. Understanding common pitfalls can save significant debugging time.

1. **Incorrect Cron Expressions:** A very common mistake is misconfiguring cron expressions, leading to jobs not running or running at unexpected times.
    - **Symptom:** Your scheduled job either doesn't run at all, runs at the wrong frequency, or at an unexpected time.
    - **Troubleshooting:** Use an online cron expression validator (e.g., crontab.guru) to test and visualize your expressions. Be mindful of time zones; Trigger.dev schedules are typically evaluated in UTC. If your local machine or dashboard displays times in a different zone, there might be a perceived offset.
2. **Idempotency Issues with Retries:** If a job step is retried, it might execute multiple times. If the operation isn't designed to be idempotent, this can lead to unintended side effects.
    - **Symptom:** Duplicate data entries, multiple notifications sent, or inconsistent state in external systems after failures and retries.
    - **Troubleshooting:** Design your job steps to be **idempotent**. This means that performing the operation multiple times has the same effect as performing it once. For example, when updating a database record, use `UPSERT` (update or insert) instead of just `INSERT`. When sending notifications, ensure your notification system can handle duplicate requests gracefully or use a unique transaction ID to prevent re-sends.
3. **Debugging Long-Running Workflows:** Understanding the exact state of a job that pauses for minutes or hours can be tricky without proper logging.
    - **Symptom:** A job appears "stuck," or you can't tell which specific step it's currently on, especially if it's waiting for external input or a long `io.wait`.
    - **Troubleshooting:** Leverage `io.logger.info` and `io.logger.error` extensively. These logs are persisted with the job run in the Trigger.dev dashboard, giving you a clear, step-by-step timeline of execution and the exact state at each durable step. The Trigger.dev dashboard provides a visual trace of your job's execution, showing each `io` call and its status, which is invaluable for pinpointing issues.

4. **Resource Exhaustion for Highly Concurrent Queued Jobs:** While queues prevent immediate overload by buffering tasks, a massive backlog can still consume significant resources over time.
    - **Symptom:** Workers become slow, experience memory issues, or jobs take an excessively long time to process, even if they eventually complete.
    - **Troubleshooting:** Monitor your queue depth and worker resource usage (CPU, memory). If your backlog consistently grows or workers are constrained, scale your Trigger.dev workers horizontally. For extremely high-throughput scenarios, consider partitioning your events (e.g., using different event names or `queue` properties if Trigger.dev introduces them for finer control) to distribute load across more workers or queues.

Summary

You've now mastered the foundational elements for building robust and reliable workflows with Trigger.dev:

- **Queues:** Decouple task producers from consumers, enabling load leveling and graceful handling of processing spikes. Every Trigger.dev job leverages an underlying managed queue for resilience and scalability.
- **Scheduling:** Execute tasks at precise times or intervals using familiar Cron expressions, perfect for reports, maintenance, or recurring checks, eliminating the need for external schedulers.
- **Durable Execution:** Trigger.dev's ability to persist job state ensures that long-running workflows, pauses (`io.wait`), and retries (`io.runTask`) can survive worker restarts and continue exactly where they last left off, guaranteeing progress.
- **Retries:** Built-in, configurable retry mechanisms handle transient failures in external API calls or other operations, making your jobs resilient to temporary service outages without complex manual coding.

These capabilities are indispensable for any production system, especially those orchestrating complex AI agents, data pipelines, or critical business logic. By understanding and applying these principles, you can build systems that are not only powerful but also fault-tolerant and highly available.

In the next chapter, we'll take these robust workflow capabilities and apply them to the exciting world of AI agents, exploring how Trigger.dev can manage the lifecycle and interactions of intelligent systems, leveraging its durable execution for complex, multi-step agentic behaviors.

References

- Trigger.dev Documentation
- Trigger.dev Jobs Reference
- Trigger.dev Event Triggers
- Trigger.dev Cron Triggers
- Trigger.dev io.wait API
- Trigger.dev io.runTask API
- Crontab Guru - Cron Schedule Expression Editor

This page is AI-assisted and reviewed. It references official documentation and recognized resources where relevant.

CHAPTER 05

Integrating Trigger.dev: Next.js, TypeScript, and External APIs

Integrating powerful backend services with a dynamic frontend is a cornerstone of modern web application development. In this chapter, we'll connect the dots between your Next.js application, robust TypeScript-powered Trigger.dev workflows, and external APIs. This combination allows you to offload heavy computations, long-running tasks, and complex integrations to Trigger.dev, keeping your frontend responsive and scalable.

You'll learn how to invoke Trigger.dev jobs from your Next.js application, define these jobs with the clarity and safety of TypeScript, and make secure, resilient calls to third-party APIs from within your workflows. This approach is fundamental for building features like automated data synchronization, background processing, and AI-driven responses without blocking your user interface. We'll be using Trigger.dev v4-beta, which is the latest iteration, with v3 being the current stable release. Version 4 is expected to go GA around May/June 2026.

Before diving in, ensure you have a basic understanding of Next.js fundamentals, including API routes, and have completed the initial Trigger.dev project setup from previous chapters. We'll be building upon that foundation to create a truly integrated system.

Trigger.dev in a Modern Web Stack

When building applications with frameworks like Next.js, you often encounter tasks that are too long-running, too resource-intensive, or too sensitive to execute directly on the frontend or within a quick API response. This is where Trigger.dev shines, acting as a reliable orchestrator for these background operations.

The Frontend-Backend Workflow Bridge

Imagine a user submits a form that requires calling multiple external APIs, processing data, and then sending an email. Doing all of this synchronously in a Next.js API route could lead to timeouts, poor user experience, and difficult error handling.

This is where Trigger.dev steps in. Your Next.js application (specifically, an API route) can act as a lightweight trigger, initiating a Trigger.dev workflow. The workflow then handles all the heavy lifting in the background, allowing your Next.js API route to respond quickly to the user.

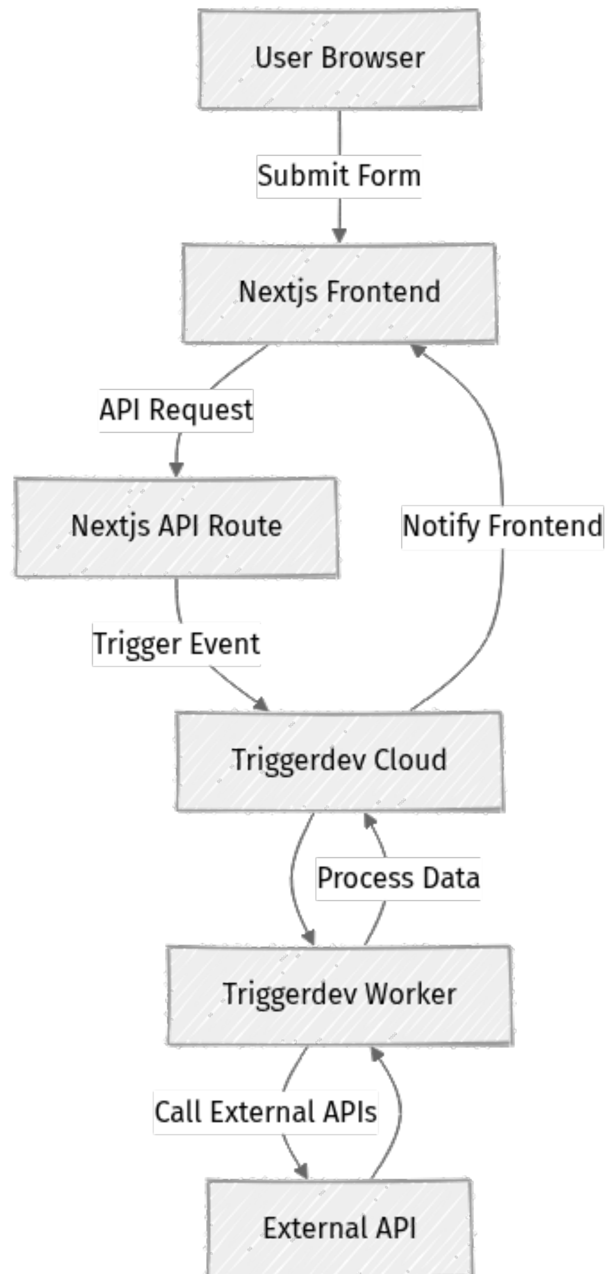

📌 Key Idea: Trigger.dev decouples the request from heavy processing, making your frontend responsive and your backend resilient. It handles the complexity of background jobs, retries, and durable execution.

Why TypeScript for Trigger.dev Workflows?

TypeScript brings static typing to JavaScript, offering significant advantages, especially for complex and long-running workflows that often involve diverse data structures and external integrations:

- **Type Safety at Compile Time:** Define the expected structure of your job payloads, outputs, and any data passed between steps. This catches many common programming errors related to data mismatches before your code even runs, preventing runtime surprises.
- **Enhanced Developer Experience:** Your Integrated Development Environment (IDE) can provide intelligent autocompletion, real-time error checking, and assist with refactoring. This makes development faster, more confident, and less prone to errors.
- **Improved Code Readability and Maintainability:** Clearly defined types act as a form of self-documentation. It becomes easier for you and your team to understand what data a workflow expects, what it produces, and how different parts of the system interact.
- **Reduced Bugs in Production:** By enforcing strict type contracts, TypeScript helps prevent unexpected data formats or missing properties from causing issues deeper within your workflow or when interacting with external systems. This is particularly crucial for long-running processes where debugging can be more challenging.

For production systems, TypeScript is almost a necessity for maintaining code quality, predictability, and team collaboration.

Securely Interacting with External APIs

Trigger.dev workflows are an ideal place to interact with external APIs because they run in a secure, server-side environment. This allows you to:

- **Protect Sensitive Credentials:** Store API keys, tokens, and other sensitive credentials securely as environment variables or using a secrets manager. These are never exposed to the client-side, mitigating security risks.
- **Implement Robust Retries and Backoff:** External APIs can be flaky. Trigger.dev's durable execution automatically retries API calls that fail due to transient network issues, timeouts, or temporary service unavailability. You can configure exponential backoff strategies to prevent overwhelming the external service.

- **Manage API Rate Limits:** Implement logic to respect API rate limits, potentially pausing or delaying subsequent calls using Trigger.dev's built-in delay functions, without impacting your frontend's responsiveness or blocking other operations.
- **Transform and Validate Data:** Process, validate, and transform data received from external APIs before it's used elsewhere or stored. This ensures data consistency and integrity within your application.

⚡ Real-world insight: Never expose API keys directly in client-side code. Always route API calls through a secure backend or a service like Trigger.dev to protect your credentials and manage API interactions robustly.

Step-by-Step Implementation

Let's build a simple Next.js application that triggers a Trigger.dev job. This job will then call a public external API and return a processed result.

Prerequisites

Before starting this section, ensure you have a Trigger.dev project set up and running, as discussed in previous chapters. This typically involves a separate Node.js project where your Trigger.dev jobs are defined and run by the Trigger.dev CLI.

1. Set Up a Next.js Project

If you don't have one already, create a new Next.js project. We'll use the App Router for modern Next.js development, along with TypeScript.

```
npx create-next-app@latest my-trigger-app --typescript --app --tailwind --
eslint
```

When prompted, choose your preferred options. For simplicity, we'll stick with the defaults. Once created, navigate into your new project directory:

```
cd my-trigger-app
```

2. Install Trigger.dev Client in Next.js

Inside your Next.js project, install the `@trigger.dev/sdk` client library. This package allows your frontend or API routes to communicate with your Trigger.dev workflows by sending events.

```
npm install @trigger.dev/sdk
# or yarn add @trigger.dev/sdk
```

3. Integrate Trigger.dev Client and Create an API Route

We'll create a Next.js API route that receives a request from the frontend and then triggers a Trigger.dev job.

First, create a `triggerClient.ts` file to initialize your Trigger.dev client. This ensures you only initialize it once and provides a single point of configuration.

Create `src/lib/triggerClient.ts` within your Next.js project:

```
// src/lib/triggerClient.ts
import { TriggerClient } from "@trigger.dev/sdk";

// Initialize the Trigger.dev client for sending events.
// The 'id' should be a unique identifier for your application.
// The 'apiKey' is your Trigger.dev API key, kept secret.
// 'apiUrl' is optional, defaults to the public Trigger.dev API.
export const client = new TriggerClient({
  id: "my-nextjs-app", // A unique ID for your application connecting to
Trigger.dev
  apiKey: process.env.TRIGGER_API_KEY, // Your secret API key
  apiUrl: process.env.TRIGGER_PUBLIC_API_URL || "https://api.trigger.dev",
});
```

Next, create an API route that will be responsible for triggering a job. This route will receive data from your frontend and forward it as a payload to a Trigger.dev event.

Create `src/app/api/trigger-job/route.ts`:

```
// src/app/api/trigger-job/route.ts
import { NextResponse } from "next/server";
import { client } from "@/lib/triggerClient"; // Import our initialized
Trigger.dev client

// Define the job's ID - this must match the ID defined in your Trigger.dev
workflow.
const MY_JOB_ID = "process-public-api-data";

export async function POST(request: Request) {
  try {
    // Parse the incoming JSON request body
    const { message } = await request.json();

    // Basic validation for the 'message' field
    if (!message) {
      return NextResponse.json({ error: "Message is required" }, { status:
400 });
    }

    // Trigger the Trigger.dev job by sending an event.
```

```
    // The 'name' property must match the 'name' in the eventTrigger of your
job definition.
    // The 'payload' is the data you want to send to your job.
    const jobRun = await client.sendEvent({
      name: MY_JOB_ID, // The name of the event that triggers your job
      payload: { userInput: message, timestamp: new
Date().toISOString() }, // Data sent to the job
    });

    console.log(`Triggered job run: ${jobRun.id}`);

    // Respond to the frontend indicating success and providing the job run ID
    return NextResponse.json({
      success: true,
      jobRunId: jobRun.id,
      message: "Job successfully triggered!",
    });
  } catch (error) {
    console.error("Error triggering job:", error);
    // Return an error response if something goes wrong
    return NextResponse.json(
      { error: "Failed to trigger job", details: (error as Error).message },
      { status: 500 }
    );
  }
}
```

Finally, you need to add your Trigger.dev API Key to your Next.js project's environment variables. Create a `.env.local` file at the root of your Next.js project:

```
# .env.local (in your Next.js project root)
TRIGGER_API_KEY=tr_dev_YOUR_SECRET_KEY_HERE
TRIGGER_PUBLIC_API_URL=https://api.trigger.dev # Optional: defaults to the
public API
```

Important: Replace `tr_dev_YOUR_SECRET_KEY_HERE` with your actual Trigger.dev API key, which you can find in your Trigger.dev dashboard. For production deployment, ensure these environment variables are configured in your hosting provider (e.g., Vercel, Netlify).

4. Create a TypeScript Workflow in Trigger.dev

Now, let's define the actual Trigger.dev job that will be executed. This will reside in your Trigger.dev project (which you should have set up in previous chapters, likely in a `src/jobs` directory).

First, ensure you have `axios` installed in your Trigger.dev project to make HTTP requests:

```
# Navigate to your Trigger.dev project directory
cd path/to/your/trigger-dev-project
```

```
npm install axios
# or yarn add axios
```

Next, create `src/jobs/process-public-api-data.ts` in your Trigger.dev project:

```
// src/jobs/process-public-api-data.ts
import { client } from "@/trigger"; // Assuming your Trigger.dev client is
initialized here (e.g., in src/trigger.ts)
import { eventTrigger } from "@trigger.dev/sdk";
import axios from "axios"; // For making HTTP requests in the job

// Define the input type for our job payload, leveraging TypeScript for
safety.
// This type must match the 'payload' structure sent from the Next.js API
route.
interface ProcessApiPayload {
  userInput: string;
  timestamp: string;
}

// Define the output type for our job, which will be the result of the 'run'
function.
interface ProcessApiResult {
  originalInput: string;
  processedData: any; // In a real app, define a more specific type for your
API response
  externalApiUrl: string;
  jobRunId: string;
}

client.defineJob({
  id: "process-public-api-data", // This ID MUST match the MY_JOB_ID in your
Next.js API route
  name: "Process Public API Data",
  version: "1.0.0",
  enabled: true,
  // This job is triggered by an event. The 'name' property here
  // must match the 'name' in client.sendEvent() from your Next.js app.
  trigger: eventTrigger({
    name: "process-public-api-data",
    schema: {
      // Define the JSON schema for the incoming payload.
      // This provides runtime validation and TypeScript inference.
      type: "object",
      properties: {
        userInput: { type: "string", description: "User provided search
term" },
        timestamp: { type: "string", format: "date-time", description: "Time
of event trigger" },
      },
      required: ["userInput", "timestamp"],
      additionalProperties: false, // Disallow unexpected properties
    },
  }),
  // The 'run' function contains the core logic of your job.
  // 'payload' will be type-checked against ProcessApiPayload.
  // '{ logger, id: jobRunId }' provides logging and the unique ID for the
current job run.
```

```
  run: async (payload: ProcessApiPayload, { logger, id: jobRunId }) => {
    logger.info("Starting job to process public API data...", { payload });

    const externalApiUrl = "https://api.publicapis.org/entries"; // A simple
public API for demonstration

    try {
      // Step 1: Call an external API using axios.
      logger.info(`Calling external API: ${externalApiUrl}`);
      const response = await axios.get(externalApiUrl, {
        params: { title: payload.userInput }, // Use user input as a query
parameter
        timeout: 10000, // 10 seconds timeout for the API call
      });

      // Step 2: Process the API response.
      // We'll take the top 3 entries from the response for simplicity.
      const processedData = response.data.entries ? response.data.entries.slic
e(0, 3) : [];
      logger.info("Successfully fetched and processed data from external
API.", {
        count: processedData.length,
      });

      // Step 3: Return the result, adhering to our ProcessApiResult type.
      const result: ProcessApiResult = {
        originalInput: payload.userInput,
        processedData: processedData,
        externalApiUrl: externalApiUrl,
        jobRunId: jobRunId,
      };

      logger.info("Job completed successfully.", { result });
      return result; // This result will be visible in the Trigger.dev
dashboard
    } catch (error) {
      logger.error("Failed to call or process external API.", { error: (error
as Error).message });
      // Throwing an error will cause Trigger.dev to potentially retry the
job,
      // depending on your job's retry configuration.
      throw new Error(`External API call failed: ${(error as Error).message}`)
;
    }
  },
});
```

5. Create a Frontend Component to Trigger the Job

Let's create a simple form in your Next.js application that calls our API route, which in turn triggers the Trigger.dev job.

Open `src/app/page.tsx` in your Next.js project and replace its content with the following:

```
// src/app/page.tsx
"use client"; // This component needs to be a Client Component to use hooks
like useState
```

```
import { useState } from "react";

export default function Home() {
  const [message, setMessage] = useState("");
  const [response, setResponse] = useState<any>(null);
  const [loading, setLoading] = useState(false);
  const [error, setError] = useState<string | null>(null);

  // Handles the form submission
  const handleSubmit = async (e: React.FormEvent) => {
    e.preventDefault(); // Prevent default form submission behavior
    setLoading(true); // Indicate loading state
    setResponse(null); // Clear previous response
    setError(null); // Clear previous errors

    try {
      // Make a POST request to our Next.js API route
      const res = await fetch("/api/trigger-job", {
        method: "POST",
        headers: {
          "Content-Type": "application/json",
        },
        body: JSON.stringify({ message }), // Send the user's input as JSON
      });

      const data = await res.json(); // Parse the JSON response from our API
route

      // Check if the API response was not OK (e.g., status 400 or 500)
      if (!res.ok) {
        throw new Error(data.details || data.error || "Something went wrong");
      }

      setResponse(data); // Set the successful response
    } catch (err) {
      setError((err as Error).message); // Catch and display any errors
    } finally {
      setLoading(false); // End loading state
    }
  };

  return (
    <main className="flex min-h-screen flex-col items-center justify-center
p-24 bg-gray-100">
      <h1 className="text-4xl font-bold mb-8 text-gray-800">Trigger.dev with N
ext.js & APIs</h1>

      <form onSubmit={handleSubmit} className="bg-white p-8 rounded-lg shadow-
md w-full max-w-md">
        <div className="mb-4">
          <label htmlFor="message" className="block text-gray-700 text-sm
font-bold mb-2">
            Enter a search term for public APIs:
          </label>
          <input
            type="text"
            id="message"
            className="shadow appearance-none border rounded w-full py-2 px-3
text-gray-700 leading-tight focus:outline-none focus:shadow-outline"
            value={message}
            onChange={(e) => setMessage(e.target.value)}
```

```
            placeholder="e.g., animals, health, food"
            required
          />
        </div>
        <button
          type="submit"
          className="bg-blue-500 hover:bg-blue-700 text-white font-bold py-2
px-4 rounded focus:outline-none focus:shadow-outline disabled:opacity-50"
          disabled={loading} // Disable button while loading
        >
          {loading ? "Triggering..." : "Trigger Trigger.dev Job"}
        </button>
      </form>

      {loading && <p className="mt-4 text-blue-600">Loading...</p>}
      {error && <p className="mt-4 text-red-600">Error: {error}</p>}
      {response && (
        <div className="mt-8 bg-green-100 border-l-4 border-green-500 text--
green-700 p-4 w-full max-w-md" role="alert">
          <p className="font-bold">Job Triggered Successfully!</p>
          <p>Job Run ID: <code className="font-mono">{response.jobRunId}</
code></p>
          <p>Message: {response.message}</p>
          <p className="mt-2">Check your Trigger.dev dashboard for job progres
s and results!</p>
        </div>
      )}
    </main>
  );
}
```

6. Run Your Applications

Now that both your Next.js application and Trigger.dev project are configured, let's bring them to life!

1. **Start your Trigger.dev project:** Open your terminal, navigate to your Trigger.dev project directory, and start the development server. This connects your local Trigger.dev worker to the Trigger.dev cloud.

```
# In your Trigger.dev project directory
npm run dev
```

```
You should see output indicating that your Trigger.dev project is running and
listening for events from the cloud.
```

1. **Start your Next.js application:** Open a separate terminal, navigate to your Next.js project directory (`my-trigger-app`), and start the Next.js development server.

```
# In your Next.js project directory
npm run dev
```

```
Open your browser to `<http://localhost:3000`>.
```

Now, enter a search term in the Next.js form (e.g., "animals", "health", "food") and submit it. You should see a "Job Triggered Successfully!" message with a `jobRunId`. Go to your Trigger.dev dashboard (usually `<http://localhost:8080`> if self-hosting, or your cloud dashboard), and you'll see a new run for the "Process Public API Data" job, showing its progress and eventual result, including the data fetched from the external API.

Mini-Challenge: Enhance the API Call

Your turn! Let's make the external API interaction a bit more dynamic and give the user more control.

- **Challenge:** Modify the `process-public-api-data.ts` Trigger.dev job to allow the user to specify which public API to call (e.g., `<https://api.publicapis.org/entries`> or `<https://catfact.ninja/fact`>).
    - Update the `ProcessApiPayload` TypeScript interface to include an `apiEndpoint` field (e.g., `apiEndpoint: string;`).
    - Modify the `eventTrigger` schema in your Trigger.dev job to include the `apiEndpoint` property.
    - Adjust the Next.js frontend (`src/app/page.tsx`) to include a new input field (like a text input or dropdown) for the `apiEndpoint`.
    - Modify the Trigger.dev job's `run` function to use the `apiEndpoint` from the payload for its `axios.get` call.
- **Hint:** Remember to handle potential errors gracefully if the provided `apiEndpoint` is invalid or inaccessible. You might want to add a default API endpoint in your job as a fallback.
- **What to observe/learn:** How to make your Trigger.dev workflows more flexible and configurable through dynamic inputs, and how TypeScript helps maintain type safety even with changing data structures.

Common Pitfalls & Troubleshooting

Working with integrated systems like Next.js and Trigger.dev can introduce a few common challenges. Here's how to debug them effectively.

⚠️ What can go wrong: Environment Variable Mismatches

- **Issue:** Your Trigger.dev API key (`TRIGGER_API_KEY`) or other secrets are not correctly loaded, leading to `401 Unauthorized` errors when your Next.js app tries to trigger jobs, or your Trigger.dev worker cannot connect.
- **Troubleshooting:**
    - **Local Development (Next.js):** Ensure `TRIGGER_API_KEY` is present in your `.env.local` file at the root of your Next.js project. Remember to restart your Next.js dev server (`npm run dev`) after changing `.env.local`, as Next.js caches these variables.
    - **Deployment (Next.js):** Verify that your hosting provider (e.g., Vercel, Netlify) has the `TRIGGER_API_KEY` environment variable correctly configured for your Next.js application. It's often set in the dashboard settings.
    - **Trigger.dev Project:** Similarly, ensure your Trigger.dev project has its `TR_API_KEY` (or the equivalent variable you used to initialize its `TriggerClient`) correctly set for its environment, both locally and in deployment.

⚠️ What can go wrong: TypeScript Configuration Errors

- **Issue:** You encounter compilation errors related to types (e.g., "Property 'x' does not exist on type 'Y'"). This often happens when the expected data structure doesn't match the actual data.

- **Troubleshooting:**
  - **`tsconfig.json`:** Ensure your `tsconfig.json` files in both your Next.js and Trigger.dev projects are correctly configured. Pay special attention to `baseUrl` and `paths` if you're using absolute imports (like `@/lib/triggerClient`).
  - **Interface/Type Mismatches:** Double-check that the `payload` you are sending from Next.js (defined by `client.sendEvent`'s `payload`) exactly matches the `schema` defined in `eventTrigger` and the `interface` (`ProcessApiPayload`) used in your Trigger.dev job. TypeScript helps catch this early, but if you change one without the other, it will break.
  - **Missing Type Definitions:** Ensure all necessary type definition packages are installed (e.g., `@types/axios` if you're using `axios` in a TypeScript project).

⚠️ What can go wrong: External API Rate Limits or Failures

- **Issue:** Your Trigger.dev jobs are failing due to frequent `429 Too Many Requests` (rate limit) or `5xx` (server error) responses from the external API.
- **Troubleshooting:**
  - **Trigger.dev Retries (Default):** By default, Trigger.dev automatically retries jobs on unhandled exceptions. This is your first line of defense against transient failures.
  - **Explicit Retries:** For more fine-grained control, you can use Trigger.dev's retry options directly in your `run` function or `defineJob` configuration. This allows you to specify maximum attempts, backoff strategies, and specific error codes to retry.

```
        // Example of explicit retries for an API call within a job step
        // (This is an illustrative example; Trigger.dev's SDK provides
dedicated retry APIs)
        client.defineJob({
          // ...
          run: async (payload, { logger }) => {
            // ...
            // Use Trigger.dev's built-in retry mechanism for a specific step
            const apiResult = await client.retries.add(
              "call-external-api", // A unique ID for this retryable step
              {
                maxAttempts: 5,
                minTimeoutInMs: 1000, // Start with 1 second delay
                maxTimeoutInMs: 30000, // Max delay of 30 seconds
                factor: 2, // Exponential backoff (1s, 2s, 4s, 8s, 16s...)
                // Other retry options like 'randomize', 'retryOn' (specific
status codes)
```

```
                },
                async () => {
                  logger.info("Attempting external API call...");
                  const response = await axios.get(externalApiUrl);
                  return response.data;
                }
              );
              // ...
            }
          });
```

```
- **Rate Limit Headers:** Many APIs include `Retry-After` headers in their
`429` responses. You can read these headers in your job and use
`client.delays.delayUntil()` to pause the job until the specified time,
effectively respecting the API's instructions.
- **Detailed Error Logging:** Always use `logger.error()` in your job to log
detailed API error responses (e.g., status codes, error messages from the
external API). This helps understand if the problem is transient or requires
code changes.
```

Summary

This chapter has guided you through the essential process of integrating Trigger.dev into a modern web application stack. You've learned:

- **Next.js as a Trigger:** How to use Next.js API routes to securely and efficiently trigger Trigger.dev jobs, effectively offloading complex and long-running tasks from your frontend and immediate API responses.
- **TypeScript for Robustness:** The significant benefits of using TypeScript to define clear types for your job payloads and outputs, leading to more maintainable, readable, and less error-prone workflows.
- **External API Interaction:** Best practices for calling external APIs from within Trigger.dev jobs, including securely handling secrets, implementing robust retry mechanisms, and managing API rate limits.

By mastering these integrations, you're now equipped to build more dynamic, scalable, and resilient applications. The ability to delegate background tasks to a durable execution platform like Trigger.dev, combined with the type safety of TypeScript, empowers you to focus on core application logic rather than infrastructure concerns.

Next, we'll delve deeper into Trigger.dev's advanced features, exploring how to manage long-running workflows and introduce human-in-the-loop interactions for even more powerful and flexible systems.

References

- Trigger.dev Documentation
- Trigger.dev GitHub Repository
- Next.js Documentation
- TypeScript Handbook
- Axios GitHub Repository

This page is AI-assisted and reviewed. It references official documentation and recognized resources where relevant.

CHAPTER 06

Observability & Debugging: Seeing Your Workflows in Action

Imagine you've launched a complex AI agent workflow or a critical data processing pipeline. Suddenly, something goes wrong: a customer report is delayed, an AI response is off, or a scheduled task simply doesn't run. Without a clear view into your system, these issues can feel like trying to debug a black box. This is where observability and debugging become your superpowers.

In modern distributed systems, especially those involving long-running processes or AI agents, it's not enough for your code to just work. You need to know how it's working, why it might be failing, and what happened at every step of its execution. Trigger.dev provides robust tools to give you this visibility, transforming opaque workflows into transparent operations.

This chapter will equip you with the knowledge and practical skills to effectively monitor and debug your Trigger.dev workflows. We'll explore Trigger.dev's built-in dashboard, understand the lifecycle of a workflow run, interpret logs and events, and learn how to proactively identify and resolve issues. By the end, you'll be able to confidently see your workflows in action and troubleshoot any hiccups.

This guide assumes you're familiar with creating basic Trigger.dev jobs and workflows, as covered in previous chapters. We'll be using Trigger.dev v4-beta (as of 2026-05-20), which is expected to go GA around May/June 2026, building upon the foundations of v3.

Core Concepts: The Eyes and Ears of Your Workflow

Observability is about understanding the internal state of a system by examining its external outputs. For Trigger.dev, this means leveraging logs, metrics, and traces to gain insight into how your background jobs and long-running workflows are performing.

What is Observability in Trigger.dev?

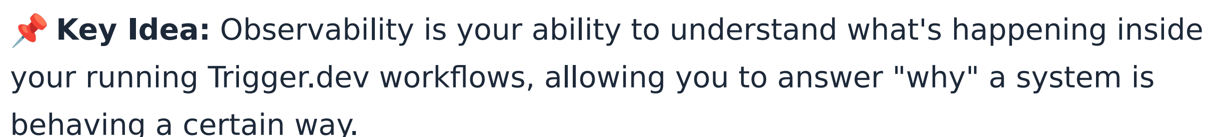

📌 **Key Idea:** Observability is your ability to understand what's happening inside your running Trigger.dev workflows, allowing you to answer "why" a system is behaving a certain way.

In the context of Trigger.dev, observability allows you to:

- **Monitor Execution:** See if your workflows are running, succeeding, failing, or retrying.
- **Diagnose Issues:** Pinpoint the exact step where an error occurred and why.
- **Track Performance:** Understand how long tasks take and identify bottlenecks.
- **Audit Actions:** Keep a record of all events and interactions within a workflow.

Why this matters: This is critical for Trigger.dev because it deals with background jobs, scheduled tasks, and durable execution. These are often asynchronous and distributed by nature. You can't just attach a debugger and step through code that might be paused for hours or retrying across different instances. You need persistent, centralized visibility to manage these complex, potentially long-running processes.

Trigger.dev's Dashboard: Your Mission Control

The Trigger.dev dashboard is your primary interface for observing and debugging your workflows. It provides a comprehensive view of all activity related to your projects. When you log into the Trigger.dev cloud, you'll find sections dedicated to:

- **Runs:** A list of every time a job has been executed.
- **Tasks:** The individual steps within a job's run.
- **Events:** The raw data that triggered or flowed through your jobs.
- **Logs:** The output from your workflow code, including `console.log` statements and error messages.

Let's visualize how these components connect during a workflow's execution.

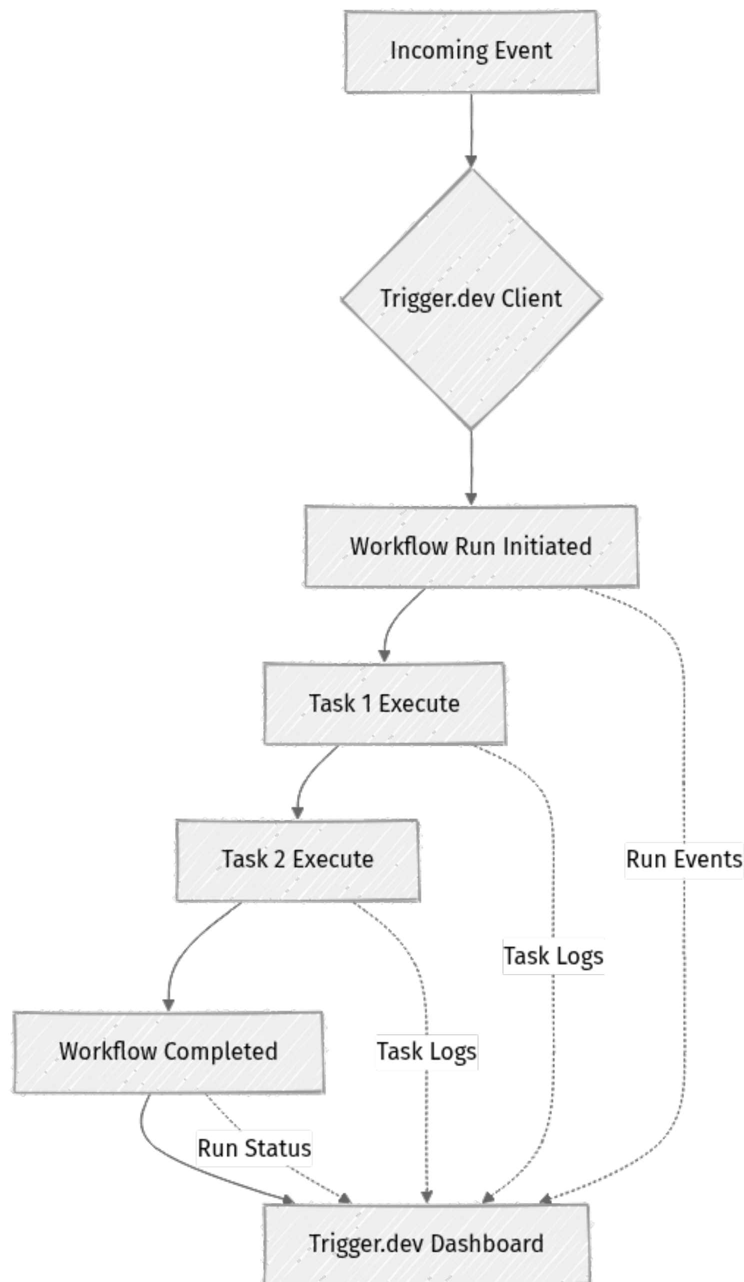

As you can see, every significant action within your workflow, from the initial event to task execution and completion, feeds information back into the Trigger.dev dashboard. This centralized view is invaluable for understanding complex, distributed flows.

Understanding Workflow Runs

A "Run" in Trigger.dev represents a single execution instance of a defined job. If your `sendWelcomeEmail` job is triggered for 10 different users, that's 10 separate runs. Each run has a unique ID and a distinct lifecycle.

Common Run Statuses:

- **Pending:** The run has been scheduled but hasn't started execution yet.

- **Running:** The workflow is currently executing one or more of its tasks.
- **Success:** All tasks in the workflow completed without unhandled errors.
- **Failed:** One or more tasks in the workflow encountered an unhandled error, and all retries were exhausted.
- **Cancelled:** The run was explicitly stopped before completion.
- **Timed Out:** The workflow exceeded its configured execution time limit.

When you click on a specific run in the dashboard, you'll get a detailed timeline of its execution, including the sequence of tasks, their individual statuses, and any associated logs. This visual timeline is incredibly helpful for quickly grasping the flow, especially for long-running processes.

Task Executions & Logs

Each step you define within a `client.defineJob` using `io.runTask` is treated as a distinct "Task" by Trigger.dev. This granular approach allows for durable execution, retries, and detailed observability. When a task executes, any `io.logger.info`, `io.logger.warn`, or `io.logger.error` statements from your code are captured and sent to the Trigger.dev dashboard. Even `console.log` statements will appear, but `io.logger` offers richer context.

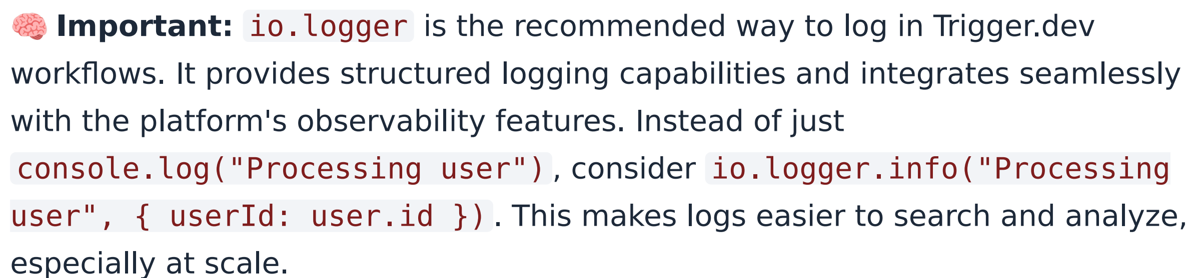

🧠 **Important:** `io.logger` is the recommended way to log in Trigger.dev workflows. It provides structured logging capabilities and integrates seamlessly with the platform's observability features. Instead of just `console.log("Processing user")`, consider `io.logger.info("Processing user", { userId: user.id })`. This makes logs easier to search and analyze, especially at scale.

Accessing individual task logs is crucial for debugging. If a workflow fails, you can often trace the error back to a specific task and then review its logs to understand the root cause. Trigger.dev also provides context around each log entry, such as the timestamp and the task ID, further aiding your investigation.

Events: The Workflow's Timeline

Events are the backbone of Trigger.dev's system. They represent significant occurrences, such as:

- An external webhook being received.
- A scheduled job being initiated.
- A task starting or completing.
- An error being thrown.

- A retry attempt.

In the Trigger.dev dashboard, the "Events" section provides a chronological feed of these activities. By examining the event stream for a particular run, you can reconstruct the exact sequence of operations, understand state transitions, and verify if external inputs were received as expected. This is particularly useful for auditing and understanding complex interactions, giving you a precise timeline of every action taken by your workflow.

Retries: Handling Transient Failures Gracefully

Trigger.dev's durable execution includes automatic retries for transient failures. This is a powerful feature, allowing your workflows to recover from temporary issues like network glitches or API rate limits. However, it also means that a task might "fail" multiple times before eventually succeeding or truly failing.

When debugging, it's vital to:

1. **Observe Retry Counts:** See how many times a task attempted to run.
2. **Understand Backoff:** Notice the increasing delays between retry attempts.
3. **Distinguish Transient vs. Permanent:** Determine if the error is something that might resolve itself (e.g., a temporary network issue) or a fundamental bug in your code that requires a fix.

The dashboard will clearly show when a task is retrying, which helps you differentiate between a single, recoverable failure and a persistent problem that needs your immediate attention.

Step-by-Step Implementation: Adding Observability to a Workflow

Let's take a simple Trigger.dev workflow and enhance its observability. We'll add custom logs and then intentionally introduce a failure to see how Trigger.dev's dashboard helps us debug.

First, ensure you have a Trigger.dev project set up. If not, you can quickly create one:

```
# As of 2026-05-20, using v4-beta
npx trigger.dev@v4-beta init
```

Follow the prompts, selecting a `Next.js` project (or your preferred framework) and giving it a name. This will create a basic project with a `trigger.ts` file in your `src/jobs` directory (or similar, depending on your framework).

Now, open your `src/jobs/trigger.ts` (or equivalent) file. We'll modify a simple job incrementally.

1. Define a Basic Job

Let's start with a minimal job definition. If you already have one, you can adapt it.

```
// src/jobs/trigger.ts
import { client } from "@/trigger";
import { eventTrigger } from "@trigger.dev/sdk";

client.defineJob({
  id: "observability-example",
  name: "Observability Example Workflow",
  version: "1.0.0",
  enabled: true,
  trigger: eventTrigger({
    name: "observability.event",
    schema: {
      type: "object",
      properties: {
        message: { type: "string" },
      },
      required: ["message"],
      additionalProperties: false,
    },
  }),
  run: async (payload, io, ctx) => {
    // This is where our workflow logic will go
    await io.runTask("initial-step", async () => {
      return `Received: ${payload.message}`;
    });
    return { status: "success", finalMessage: payload.message };
  },
});
```

This job simply defines an `initial-step` and returns a success message. It's a good starting point to add observability.

2. Add Custom Logging with `io.logger`

Now, let's enhance this job by adding more meaningful logs using `io.logger`. This will give us visibility into the workflow's progression.

Replace the `run` function in your `src/jobs/trigger.ts` with the following:

```
// src/jobs/trigger.ts (partial update to the run function)
// ... (previous imports and client.defineJob setup)
  run: async (payload, io, ctx) => {
    // 1. Log the incoming payload at the start of the run
```

```
    io.logger.info("Workflow started with payload", payload);

    await io.runTask("step-one", async () => {
      // 2. Log progress within a task
      io.logger.info("Starting step one: Processing message...");
      await new Promise((resolve) => setTimeout(resolve, 1000)); // Simulate
work
      io.logger.info(`Step one completed for message: ${payload.message}`);
      return `Processed: ${payload.message}`;
    });

    await io.runTask("step-two", async () => {
      io.logger.info("Starting step two: Finalizing workflow...");
      await new Promise((resolve) => setTimeout(resolve, 500)); // Simulate
more work
      io.logger.info("Step two completed successfully.");
      return "Successfully finalized.";
    });

    io.logger.info("Workflow finished successfully!");
    return { status: "success", finalMessage: payload.message };
  },
});
```

Here's what we've added and why:

- **`io.logger.info("Workflow started with payload", payload);`**: We're using `io.logger.info` to log the entire incoming payload. This is crucial for understanding what inputs triggered the job, right at the beginning of its execution.
- **Progress Logs within Tasks**: We add `io.logger.info` calls at the start and end of each `io.runTask` to clearly delineate its execution and progress. This is especially helpful for long-running tasks, as it confirms which part of the task is currently active.

3. Introduce an Intentional Failure

To truly test our debugging skills, let's introduce a conditional failure based on the incoming payload. This will allow us to observe a failed run in the dashboard.

First, update the `eventTrigger` schema to accept a `shouldFail` boolean property.

```
// src/jobs/trigger.ts (partial update to the eventTrigger schema)
// ...
  trigger: eventTrigger({
    name: "observability.event",
    schema: {
      type: "object",
      properties: {
        message: { type: "string" },
        shouldFail: { type: "boolean" }, // Add this property
      },
```

```
      required: ["message", "shouldFail"], // Make it required
      additionalProperties: false,
    },
  }),
// ...
```

Next, modify the `run` function again to incorporate the failure logic in `step-two`.

```
// src/jobs/trigger.ts (partial update to the run function)
// ... (previous imports and client.defineJob setup)
  run: async (payload, io, ctx) => {
    io.logger.info("Workflow started with payload", payload);

    await io.runTask("step-one", async () => {
      io.logger.info("Starting step one: Processing message...");
      await new Promise((resolve) => setTimeout(resolve, 1000)); // Simulate
work
      io.logger.info(`Step one completed for message: ${payload.message}`);
      return `Processed: ${payload.message}`;
    });

    await io.runTask("step-two", async () => {
      io.logger.info("Starting step two: Checking for failure condition...");
      // Introduce an intentional failure based on payload
      if (payload.shouldFail) {
        io.logger.error("Failure condition met! Throwing an error.");
        throw new Error("Deliberate failure as requested by payload.");
      }
      io.logger.info("Step two completed successfully. No failure.");
      return "Successfully avoided failure.";
    });

    io.logger.info("Workflow finished successfully!");
    return { status: "success", finalMessage: payload.message };
  },
});
```

Now, our `step-two` will check the `payload.shouldFail` property. If `true`, it will log an error and throw an exception, simulating a workflow failure.

4. Run and Observe

Let's see our enhanced observability in action.

1. **Start your Trigger.dev development server:** Open your terminal in your project's root directory and run:

```
npm run dev
```

```
This will connect your local project to your Trigger.dev cloud project,
allowing you to trigger jobs and see their runs.
```

1. **Trigger the workflow from the Trigger.dev Dashboard:**
    - Open your web browser and go to your Trigger.dev dashboard (e.g., `<https://cloud.trigger.dev>`).
    - Navigate to your specific project.
    - In the left sidebar, find the "Observability Example Workflow" job under the "Jobs" section.
    - Click on the job, then click the "Trigger" or "Run Now" button.
    - You'll be prompted for a payload.
2. **First Run (Success Case):**
    - Enter the following JSON payload into the input box:

```
        {
          "message": "Hello from Trigger.dev!",
          "shouldFail": false
        }
```

```
- Click "Run".
- Observe the "Runs" section. A new run should appear and quickly transition
to "Success".
```

1. **Second Run (Failure Case):**
    - Trigger the workflow again using the "Trigger" or "Run Now" button.
    - Enter the following JSON payload:

```
        {
          "message": "This run should fail.",
          "shouldFail": true
        }
```

```
- Click "Run".
- Observe the "Runs" section. This run should eventually show a "Failed"
status.
```

5. Inspecting the Dashboard: A Guided Tour

Now, let's dive into the dashboard to understand what happened in both cases.

1. **Navigate to the "Runs" tab** in your Trigger.dev project (usually the default view after selecting a job).
2. **Click on the "Success" run.**
    - You'll see a visual timeline of `step-one` and `step-two`. Both should clearly show "Success".
    - Below the timeline, you'll find tabs such as "Payload", "Logs", "Events", "Output", etc.
    - Click on the **"Logs" tab**. You should see all your `io.logger.info` messages in chronological order, including "Workflow started with payload", "Starting step one...", and "Step one completed...", and "Step two completed...". Notice how the `payload` object is nicely structured in the log entry.
    - Click on the **"Events" tab**. You'll see a chronological list of events like `JOB_STARTED`, `TASK_STARTED` (for each step), `TASK_COMPLETED`, and `JOB_COMPLETED`. This forms the detailed audit trail of your workflow.
3. **Click back to the "Runs" tab and then click on the "Failed" run.**
    - The timeline will immediately highlight `step-one` as "Success" and `step-two` as "Failed". This visually pinpoints the exact task where the problem occurred.
    - Go to the **"Logs" tab**. Here, you'll find the logs for `step-one` followed by those for `step-two`. Crucially, you'll see:
        - `Starting step two: Checking for failure condition...`
        - `Failure condition met! Throwing an error.` (Your custom `io.logger.error` message)
        - The actual JavaScript error stack trace: `Error: Deliberate failure as requested by payload.` This stack trace is invaluable for identifying the exact line of code that caused the error.
    - Go to the **"Events" tab**. You'll see events like `TASK_FAILED`, `JOB_FAILED`, and potentially `TASK_RETRIED` if retries were configured (by default, an unhandled error in a `io.runTask` without specific retry options might lead to immediate failure, but with `maxAttempts` you'd see `TASK_RETRIED` events).

⚡ **Real-world insight:** This ability to trace logs and events through a distributed workflow run is invaluable. It's how you diagnose issues in production without direct access to the server, and it forms the basis of incident response and post-mortem analysis.

Mini-Challenge: Debugging a Flaky Workflow

Let's put your new observability skills to the test with a workflow that's a bit more unpredictable.

Challenge: Random Failure Generator

Modify your `observability-example` job to introduce a random failure in `step-two`. The goal is to make it fail about 30% of the time.

1. **Remove the `shouldFail` payload property** from the `eventTrigger` schema and the `run` function's `payload` type.
    - Hint: Update the `schema`'s `properties` and `required` arrays.
2. **Modify `step-two`** to randomly throw an error using `Math.random()`.
3. **Trigger the workflow multiple times** (e.g., 5-10 times) from the dashboard using any simple message payload (e.g., `{"message": "Test"}`).
4. **Observe the runs:** How many succeeded? How many failed?
5. **For a failed run:** Identify the exact task that failed and locate the error message in the logs.
6. **For a successful run (after some failures):** Note how the logs still reflect the randomness (e.g., the `randomNumber` you might log), but the workflow ultimately completed.

```
// Hint for random failure logic within step-two's async function:
// const randomNumber = Math.random();
// io.logger.info("Random number generated", { randomNumber }); // Log for observability!
// if (randomNumber < 0.3) {
//   io.logger.error(`Random failure triggered! Value: ${randomNumber}`);
//   throw new Error("Randomly decided to fail this time.");
// }
```

What to observe/learn: You'll see how even intermittent issues can be tracked down using the task-level logs and run statuses in the Trigger.dev dashboard. This simulates real-world scenarios where external APIs might occasionally return errors, or network conditions might be unstable. The key is to use logs to understand why a specific instance failed, even if most succeed.

Common Pitfalls & Troubleshooting

Even with great observability tools, certain patterns can make debugging challenging in distributed systems.

⚠️ What can go wrong: Too Much Logging vs. Too Little

- **Too much logging:** While comprehensive logs are helpful, excessive logging can overwhelm your dashboard, make important messages hard to find, and potentially incur higher costs if logs are stored externally. Imagine millions of log lines for a simple workflow; finding the needle in that haystack is tough.
  - 🔥 **Optimization / Pro tip:** Use different log levels (`info`, `warn`, `error`, `debug`) judiciously. Only log truly critical or actionable information at `info` or `warn` levels in production. Use `debug` for detailed, development-time insights that can be disabled in production.
- **Too little logging:** Conversely, if you don't log enough, you'll stare at a "Failed" status with no context, making debugging a guessing game.
  - 🔥 **Optimization / Pro tip:** Always log key inputs, outputs of complex operations, and points where external services are called or critical decisions are made. When integrating with external APIs, log the request payload and the response (or at least status codes) for easy debugging.

⚠️ What can go wrong: Misinterpreting Retry Behavior

Trigger.dev's retries are designed for transient faults. However, if your code has a persistent bug (e.g., an incorrect API key, a logic error), retries will only delay the inevitable failure and consume unnecessary resources.

- **Debugging:** If a task consistently fails and retries, check the logs for the same error message appearing repeatedly. This indicates a non-transient bug that needs a code fix, not just more retries. Don't be fooled by multiple failures; look for the pattern of failure.

- **Configuration:** Be mindful of `maxAttempts` and `retryDelay` settings. An infinite retry loop for a hard error can cause resource exhaustion and hide the true problem for longer.

⚠️ What can go wrong: State Management in Long-Running Workflows

Trigger.dev workflows are durable, meaning they can pause and resume. However, if you rely on in-memory state that isn't explicitly passed between `io.runTask` calls or stored durably, that state can be lost during pauses or retries.

- **Pitfall:** Assuming a variable set before an `await io.runTask` will retain its value if the task is retried or the workflow pauses and resumes on a different server instance.
- **Solution:** Pass necessary state as arguments to `io.runTask` or store it in a durable external system (like a database or a key-value store) if it needs to persist across long periods or multiple job runs. Each `io.runTask` should ideally be idempotent and receive all necessary context as arguments, minimizing reliance on implicit shared state.

⚠️ What can go wrong: Debugging Across Distributed Services

Your Trigger.dev workflow might interact with other microservices, databases, or external APIs. An error reported by Trigger.dev might originate from one of these external systems, not directly from your workflow code.

- **Challenge:** The Trigger.dev dashboard shows its view of the error, but the root cause might be in a different service's logs.
- **Solution:** Integrate Trigger.dev's observability with your broader observability stack (e.g., a centralized logging system like Datadog, Splunk, or Elastic Stack). Use **correlation IDs** (often passed in HTTP headers for external API calls) to link logs from Trigger.dev to logs in your other services, allowing you to trace requests end-to-end across your entire system. This provides a holistic view when debugging complex interactions.

Summary

Observability and debugging are non-negotiable skills for building robust production systems with Trigger.dev. You've learned how to:

- **Leverage the Trigger.dev Dashboard** as your central hub for monitoring workflow runs, tasks, events, and logs.

- **Understand Workflow Run Statuses** and their significance in the lifecycle of a job.
- **Interpret Task-level Logs** to pinpoint errors and track progress within individual steps, using `io.logger` for enhanced context.
- **Utilize Events** to reconstruct the exact sequence of actions and state changes, providing a detailed audit trail.
- **Recognize and Diagnose Retry Behavior** to distinguish between transient and persistent failures.
- **Implement Custom Logging** using `io.logger` to enhance visibility into your code.
- **Identify and Avoid Common Pitfalls** related to logging volume, misinterpreting retries, state management in durable workflows, and debugging across distributed services.

By mastering these concepts, you're not just building workflows; you're building reliable, transparent, and maintainable automated systems. In the next chapter, we'll shift our focus from observing individual runs to deploying your Trigger.dev applications to production environments, scaling them, and implementing best practices for real-world usage.

References

- Trigger.dev Documentation: Observability
- Trigger.dev Documentation: Logging
- Trigger.dev Documentation: Error Handling & Retries
- Trigger.dev GitHub Repository

This page is AI-assisted and reviewed. It references official documentation and recognized resources where relevant.

CHAPTER 07

Unleashing AI Agents: Building Smart, Automated Systems

Introduction

Welcome to Chapter 7! In the rapidly evolving world of software, AI agents are becoming indispensable for automating complex, multi-step tasks that require reasoning, planning, and interaction with external tools. Imagine a system that can understand a user's request, break it down into smaller problems, use various tools (like APIs or databases) to gather information, and then formulate a coherent response or take action—all without constant human supervision. That's the power of AI agents.

This chapter will guide you through building such intelligent, automated systems using Trigger.dev. We'll explore how Trigger.dev's robust features like durable execution, retries, and observability provide the perfect foundation for creating reliable and scalable AI agents. By the end, you'll understand the core concepts of agentic workflows and have built your first Trigger.dev-powered agent.

You should be familiar with the basics of Trigger.dev jobs, durable execution, and project setup from previous chapters. We'll build upon that knowledge to integrate Large Language Models (LLMs) and external tools into sophisticated workflows.

Core Concepts of AI Agents in Trigger.dev

Before we dive into code, let's establish a clear understanding of what AI agents are and why Trigger.dev is an excellent choice for orchestrating them.

What is an AI Agent?

At its heart, an AI agent is a system designed to perceive its environment, make decisions, and take actions to achieve specific goals. Unlike simple scripts that follow a predefined path, agents can adapt to new information, recover from errors, and often learn from their experiences.

A common pattern for AI agents, especially those powered by LLMs, involves a loop:

1. **Observe:** Understand the current state, user input, or system events.

2. **Plan:** Based on the observation and goal, decide the next step. This might involve choosing a tool, generating a sub-task, or refining the objective.
3. **Act:** Execute the planned step, often by calling an external tool (e.g., an API, database query, or another Trigger.dev job).
4. **Reflect:** Evaluate the outcome of the action. Did it succeed? Is the goal closer? What went wrong? Use this to refine the plan or adjust future actions.

This iterative process allows agents to tackle complex problems by breaking them down into manageable steps.

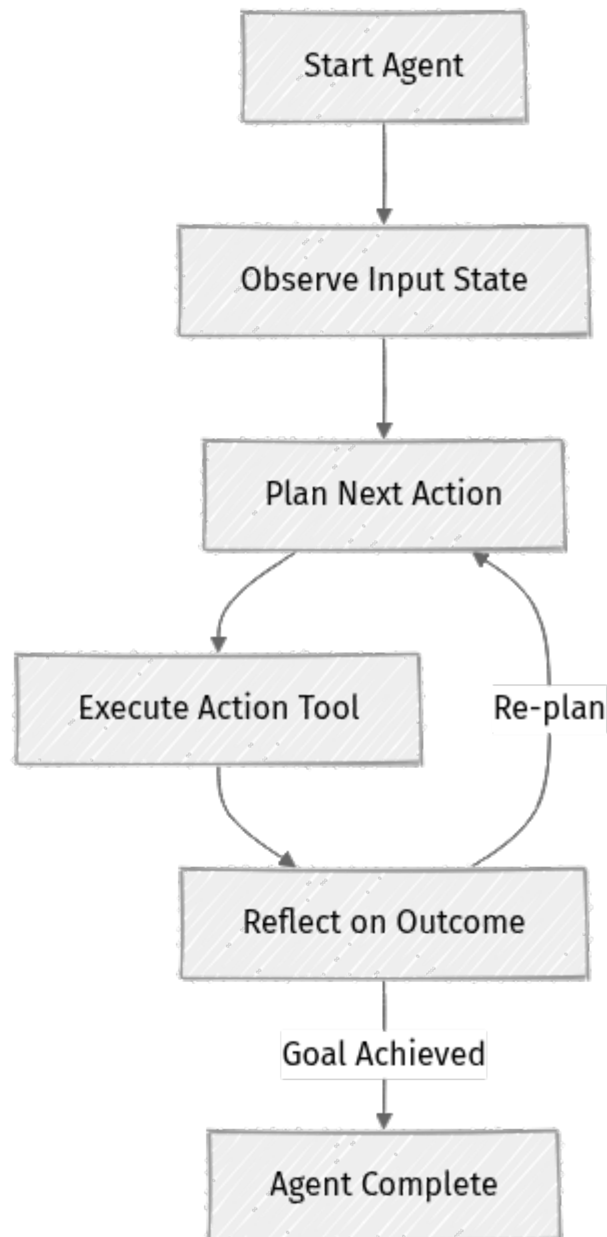

Anatomy of a basic AI Agent's decision loop.

Why Trigger.dev for Agents?

Building robust AI agents presents unique challenges:

- **Long-running tasks:** Agent workflows can involve multiple steps, API calls, and human interactions, potentially spanning minutes, hours, or even days.
- **External dependencies:** Agents rely heavily on external tools (LLM providers, third-party APIs) which can be unreliable.
- **State management:** Maintaining context and progress across multiple steps and retries is crucial.
- **Observability:** Understanding an agent's decision-making process and execution path is vital for debugging and auditing.

Trigger.dev directly addresses these challenges:

- **Durable Execution:** Trigger.dev jobs are inherently durable. If your server crashes or an external API takes too long, Trigger.dev automatically saves the agent's state and resumes it exactly where it left off. This is critical for long-running agentic workflows.
- **Built-in Retries:** When an agent attempts to use a tool (e.g., makes an API call) and it fails due to a transient error, Trigger.dev's declarative retry mechanisms automatically re-attempt the operation, improving agent resilience.
- **Queues and Concurrency:** Agents can generate many sub-tasks or calls to external tools. Trigger.dev's queuing system handles these efficiently, allowing you to scale your agent's workload without overwhelming your backend or external services.
- **Observability:** Every step of a Trigger.dev job, including tool invocations and state changes within your agent, is logged and visible in the Trigger.dev dashboard. This provides an invaluable "trace" of your agent's reasoning and actions.

⚡ **Real-world insight:** Imagine an AI agent tasked with processing customer support tickets. It might call a sentiment analysis API, then a CRM API, then draft an email, and finally wait for human approval. Each of these steps could fail or take time. Trigger.dev ensures that this complex, multi-stage process doesn't fall apart, even if the underlying infrastructure is flaky.

Agentic Workflow Design Principles

When designing agents with Trigger.dev, consider these principles:

1. **Tool-Use Focus:** Agents gain their power by interacting with tools. Design your tools as distinct, reusable functions that the LLM can invoke.
2. **Explicit State Management:** While Trigger.dev handles job state, your agent's internal reasoning state (e.g., "current plan," "remaining tasks") should be managed explicitly within your job's variables.
3. **Clear Objectives:** Define what success looks like for your agent. This helps the LLM focus and allows for better reflection.
4. **Human-in-the-Loop:** For critical or uncertain tasks, design points where a human can review, approve, or override an agent's decision. Trigger.dev's ability to pause and resume jobs is perfect for this.

MCP Integration: A Note

The term "MCP integration" is not widely documented within the public Trigger.dev ecosystem as of 2026-05-20. It's possible this refers to an internal Trigger.dev component, a specific partner integration not yet publicly detailed, or a general pattern for "Multi-Cloud Platform" or "Management Control Plane" integration.

For the purpose of this guide, we will focus on the well-established patterns of integrating AI agents with external APIs and services, which Trigger.dev excels at. If "MCP" refers to a specific Trigger.dev feature that emerges with v4 GA, it would likely involve enhanced control plane functionalities for managing complex, distributed agent deployments. Until then, treat your external LLM providers and custom tools as the primary integrations for your agents.

Building Your First Trigger.dev AI Agent

Let's get practical! We'll create a simple AI agent that takes a topic, uses an LLM to generate a short summary, and then (as a tool) "publishes" that summary.

Setting Up the Project

First, ensure you have your Trigger.dev v4-beta project set up. If you haven't yet, create a new project:

```
npx trigger.dev@v4-beta init
```

Follow the prompts to set up your project. This will create a basic Trigger.dev project with a `trigger.ts` file and an example job.

Next, we'll need an LLM client. For this example, we'll use OpenAI's API. Install the required package:

```
npm install openai@4.x.x
# or yarn add openai@4.x.x
```

(Note: As of 2026-05-20, OpenAI's official *`openai`* package is at version 4.x.x. Always check their official documentation for the latest recommended version.)

You'll also need an `OPENAI_API_KEY`. Add this to your `.env` file:

```
# .env
OPENAI_API_KEY="sk-your-openai-api-key"
```

Defining a Simple Agentic Job

We'll start by creating a job that orchestrates a simple LLM call. Open `src/trigger.ts` (or `src/jobs/myAgentJob.ts` if you prefer modular files) and add the following:

```
// src/trigger.ts
import { TriggerClient, eventTrigger } from "@trigger.dev/sdk";
import OpenAI from "openai";

// Ensure you have your Trigger.dev project ID and API key configured
export const client = new TriggerClient({
  id: "your-project-id", // Replace with your Trigger.dev Project ID
  apiKey: process.env.TRIGGER_API_KEY,
});

// Initialize OpenAI client
const openai = new OpenAI({
  apiKey: process.env.OPENAI_API_KEY,
});

client.defineJob({
  id: "ai-summarizer-agent",
  name: "AI Summarizer Agent",
  version: "1.0.0",
  trigger: eventTrigger({
    name: "summarize.topic",
    schema: {
      type: "object",
      properties: {
        topic: { type: "string" },
        length: { type: "string", default: "short" },
      },
      required: ["topic"],
      additionalProperties: false,
    }
```

```
  }),
  run: async (payload, io, ctx) => {
    await io.logger.info(`Starting summarization for topic: ${payload.topic}`)
;

    // Step 1: Call the LLM to generate a summary
    const completion = await io.runTask(
      "generate-summary-with-llm",
      async () => {
        const response = await openai.chat.completions.create({
          model: "gpt-4o", // Using the latest model as of 2026-05-20
          messages: [
            {
              role: "system",
              content: `You are an expert summarization agent. Summarize the
provided topic concisely. The summary should be ${payload.length}.`,
            },
            { role: "user", content: payload.topic },
          ],
          temperature: 0.7,
          max_tokens: 150,
        });
        return response.choices[0].message.content;
      },
      {
        name: "Generate Summary with OpenAI",
        properties: {
          topic: payload.topic,
          model: "gpt-4o",
        },
      }
    );

    if (!completion) {
      await io.logger.error("LLM did not return a summary.");
      throw new Error("Failed to generate summary.");
    }

    await io.logger.info(`Generated summary: ${completion}`);

    // Step 2: "Publish" the summary (simulate a tool call)
    await io.runTask(
      "publish-summary",
      async () => {
        // In a real agent, this would be an API call to a publishing service,
        // a database write, or a notification.
        await new Promise(resolve => setTimeout(resolve, 1000)); // Simulate
async work
        await io.logger.info(`Summary "${completion.substring(0, 30)}..."
published.`);
        return { status: "success", summary: completion };
      },
      {
        name: "Publish Summary Tool",
        properties: {
          summaryPreview: completion.substring(0, 50),
        },
      }
    );

    return {
      status: "success",
```

```
      originalTopic: payload.topic,
      generatedSummary: completion,
    };
  },
});
```

Here's a breakdown of what we added:

- **`openai` import and initialization:** We bring in the OpenAI client and initialize it using your `OPENAI_API_KEY`.
- **`client.defineJob`:** We define a new job named `ai-summarizer-agent`.
- **`eventTrigger`:** This job is triggered by a custom event `summarize.topic` which expects a `topic` (string) and optionally a `length` (string).
- **`io.logger.info`:** We use Trigger.dev's `io.logger` to log the agent's progress, which will appear in your Trigger.dev dashboard.
- **`io.runTask("generate-summary-with-llm", ...)`:** This is crucial. `io.runTask` wraps asynchronous operations, making them durable and retryable.
  - Inside, we call `openai.chat.completions.create` using the `gpt-4o` model (latest as of 2026-05-20) to generate a summary based on the `payload.topic` and `payload.length`.
  - The `name` and `properties` arguments to `io.runTask` provide excellent observability, allowing you to see exactly what the LLM call was about in the dashboard.
- **Error Handling:** We check if `completion` is returned and throw an error if not, demonstrating how Trigger.dev jobs can handle failures.
- **`io.runTask("publish-summary", ...)`:** This simulates an external tool call. In a real application, this would be an actual API call to save the summary, send it to another service, or update a database.
- **Return Value:** The job returns a structured object indicating success and providing the original topic and generated summary.

To run this locally, ensure your Trigger.dev development server is active:

```
npm run dev
```

Then, trigger the event from your Trigger.dev dashboard or using the `io.sendEvent` method in another job. For example, you can send an event via the dashboard's "Send an event" feature with:

```
{
  "name": "summarize.topic",
  "payload": {
    "topic": "The history of quantum computing",
    "length": "medium"
  }
}
```

Observe how Trigger.dev logs each step, including the LLM call and the simulated publish action, providing a clear trace of your agent's execution.

Adding Tools: Enhancing Agent Capabilities

Our current agent is simple. True AI agents leverage **tools**—functions that allow them to interact with the external world beyond just generating text. Let's imagine our agent needs to fetch information before summarizing.

We'll add a simple "search" tool. For demonstration, this tool will just return a predefined string, but in a real scenario, it would call a search API (e.g., Google Search, internal knowledge base).

First, let's define our tools. We'll add a new utility file for this, say `src/tools.ts`:

```
// src/tools.ts
import { io } from "@trigger.dev/sdk";

export async function searchTool(query: string): Promise<string> {
  return io.runTask(
    `search-tool-${query.substring(0, 20).replace(/\s/g, "-")}
`, // Dynamic ID for observability
    async () => {
      await io.logger.info(`Simulating search for: "${query}"`);
      // In a real scenario, this would be an API call to a search engine
      // e.g., const response = await axios.get(`https://api.search.com?q=$
{query}`);
      // return response.data;
      await new Promise(resolve => setTimeout(resolve, 2000)); // Simulate
network delay
      if (query.toLowerCase().includes("quantum computing")) {
        return "Quantum computing uses quantum-mechanical phenomena like
superposition and entanglement to perform computations. It can solve problems
intractable for classical computers.";
      } else if (query.toLowerCase().includes("machine learning")) {
        return "Machine learning is a subset of AI that enables systems to
learn from data, identify patterns, and make decisions with minimal human
intervention.";
      }
      return `No specific information found for "${query}".`;
    },
    {
      name: `Search Tool: ${query}`,
      properties: {
        query: query,
      },
    }
```

```
  );
}

export async function publishSummaryTool(summary: string): Promise<{status: st
ring, summary: string}> {
  return io.runTask(
    "publish-summary-tool",
    async () => {
      await io.logger.info(`Attempting to publish summary: "${summary.substrin
g(0, 50)}..."`);
      await new Promise(resolve => setTimeout(resolve, 1000)); // Simulate
publishing delay
      return { status: "success", summary: summary };
    },
    {
      name: "Publish Summary Tool",
      properties: {
        summaryPreview: summary.substring(0, 50),
      },
    }
  );
}
```

Now, let's modify our `ai-summarizer-agent` job in `src/trigger.ts` to use these tools by having the LLM decide when to call them. This involves using OpenAI's "function calling" capability.

```
// src/trigger.ts
import { TriggerClient, eventTrigger } from "@trigger.dev/sdk";
import OpenAI from "openai";
import { searchTool, publishSummaryTool } from "./tools"; // Import our new
tools

export const client = new TriggerClient({
  id: "your-project-id", // Replace with your Trigger.dev Project ID
  apiKey: process.env.TRIGGER_API_KEY,
});

const openai = new OpenAI({
  apiKey: process.env.OPENAI_API_KEY,
});

client.defineJob({
  id: "ai-smart-agent-with-tools", // New job ID
  name: "AI Smart Agent with Tools",
  version: "1.0.0",
  trigger: eventTrigger({
    name: "agent.request",
    schema: {
      type: "object",
      properties: {
        user_prompt: { type: "string" },
      },
      required: ["user_prompt"],
      additionalProperties: false,
    }
  }),
  run: async (payload, io, ctx) => {
```

```
    await io.logger.info(`Agent received request: ${payload.user_prompt}`);

    let messages: OpenAI.Chat.Completions.ChatCompletionMessageParam[] = [
      {
        role: "system",
        content: `You are a helpful AI assistant. You have access to tools to
help you answer questions or perform tasks.
                  You should first consider if you need to search for
information before answering.
                  If you generate a summary, you should use the
'publishSummaryTool' to make it available.`,
      },
      { role: "user", content: payload.user_prompt },
    ];

    const availableTools = {
      searchTool: searchTool,
      publishSummaryTool: publishSummaryTool,
    };

    const tools: OpenAI.Chat.Completions.ChatCompletionTool[] = [
      {
        type: "function",
        function: {
          name: "searchTool",
          description: "Searches for information on a given query.",
          parameters: {
            type: "object",
            properties: {
              query: {
                type: "string",
                description: "The search query.",
              },
            },
            required: ["query"],
          },
        },
      },
      {
        type: "function",
        function: {
          name: "publishSummaryTool",
          description: "Publishes a generated summary.",
          parameters: {
            type: "object",
            properties: {
              summary: {
                type: "string",
                description: "The summary text to publish.",
              },
            },
            required: ["summary"],
          },
        },
      },
    ];

    let loopCount = 0;
    const MAX_LOOP_COUNT = 5; // Prevent infinite loops

    while (loopCount < MAX_LOOP_COUNT) {
      const chatCompletion = await io.runTask(
```

```
        `agent-llm-call-${loopCount}`,
        async () => {
          return openai.chat.completions.create({
            model: "gpt-4o",
            messages: messages,
            tools: tools,
            tool_choice: "auto", // Allow the model to choose to call a tool
or respond
            temperature: 0.7,
          });
        },
        {
          name: `LLM Call ${loopCount}`,
          properties: {
            messages: messages.map(m => ({ role: m.role, content: typeof m.con
tent === 'string' ? m.content.substring(0, 100) : 'complex content' })),
            tools: tools.map(t => t.function.name),
          },
        }
      );

      const responseMessage = chatCompletion.choices[0].message;
      messages.push(responseMessage); // Add LLM's response to conversation
history

      if (responseMessage.tool_calls) {
        // Step: Act - LLM wants to call a tool
        await io.logger.info(`LLM requested tool calls: ${JSON.stringify(respo
nseMessage.tool_calls)}`);

        for (const toolCall of responseMessage.tool_calls) {
          const functionName = toolCall.function.name;
          const functionArgs = JSON.parse(toolCall.function.arguments);

          if (functionName in availableTools) {
            const toolToCall = availableTools[functionName as keyof typeof ava
ilableTools];
            const toolOutput = await toolToCall(functionArgs.query || function
Args.summary); // Pass args dynamically

            messages.push({
              tool_call_id: toolCall.id,
              role: "tool",
              name: functionName,
              content: JSON.stringify(toolOutput),
            });
            await io.logger.info(`Tool '${functionName}' executed. Output: ${J
SON.stringify(toolOutput).substring(0, 100)}`);
          } else {
            const errorMessage = `Tool '${functionName}' not found.`;
            messages.push({
              tool_call_id: toolCall.id,
              role: "tool",
              name: functionName,
              content: errorMessage,
            });
            await io.logger.error(errorMessage);
          }
        }
      } else {
        // Step: Respond - LLM has a final answer or message
        await io.logger.info(`LLM final response: ${responseMessage.content}`)
```

```
;
        return {
          status: "completed",
          finalResponse: responseMessage.content,
          conversationHistory: messages,
        };
      }
      loopCount++;
    }

    // If we reach here, the agent hit the loop limit
    await io.logger.warn("Agent reached maximum loop count without a final
response.");
    return {
      status: "failed",
      finalResponse: "Agent failed to complete task within loop limit.",
      conversationHistory: messages,
    };
  },
});
```

Let's break down the changes:

- **`ai-smart-agent-with-tools` job:** This new job takes a `user_prompt`.
- **`messages` array:** This array holds the entire conversation history, including system instructions, user prompts, LLM responses, and tool outputs. This is how the LLM maintains context.
- **`availableTools` and `tools` arrays:**
  - `availableTools` is a JavaScript object mapping tool names to their actual function implementations.
  - `tools` is an array of objects describing the tools to the OpenAI API, including their `name`, `description`, and `parameters` schema. This allows the LLM to understand when and how to call a tool.

- **Agent Loop (`while (loopCount < MAX_LOOP_COUNT)`):** This is the core of our agent.
    - **`openai.chat.completions.create` with `tools`:** The LLM is now invoked with the `tools` definition. It can decide to either generate a text response or request to call one of the provided tools. `tool_choice: "auto"` allows this flexibility.
    - **`responseMessage.tool_calls` check:** If the LLM decides to call a tool, `responseMessage.tool_calls` will be present.
    - **Tool Execution:** We iterate through the requested tool calls, parse their arguments, and then execute the actual JavaScript function from our `availableTools` object.
    - **Adding Tool Output to `messages`:** Crucially, the output of the tool call is then added back to the `messages` array with `role: "tool"`. This feeds the tool's result back to the LLM for its next turn of reasoning.
    - **Final Response:** If the LLM doesn't request a tool call, it means it has a final text response, and the job completes.
- **`MAX_LOOP_COUNT`:** A safety mechanism to prevent runaway agents.

To test this, send an event to `agent.request` (via Trigger.dev dashboard or `io.sendEvent`):

```
{
  "name": "agent.request",
  "payload": {
    "user_prompt": "Tell me about machine learning and then publish a short summary of it."
  }
}
```

Or try:```json { "name": "agent.request", "payload": { "user_prompt": "Explain quantum computing and publish a summary." } }

```
Observe how Trigger.dev logs each LLM call and each tool invocation. You'll
see the agent first calls `searchTool`, then uses that information to generate
a summary, and finally calls `publishSummaryTool`. This demonstrates the
Observe -> Plan -> Act -> Reflect loop in action, orchestrated durably by
Trigger.dev.

## Mini-Challenge: Enhance Your Agent with Human-in-the-Loop

Your challenge is to add a human approval step before the `publishSummaryTool`
is called. If the human rejects the summary, the agent should try to
regenerate it.
```

```
**Challenge:**
Modify the `ai-smart-agent-with-tools` job to include a human approval step.
1.  After the LLM generates a summary (but before `publishSummaryTool` is
called), use `io.waitforExternalService` (or a similar mechanism, perhaps by
sending an event to a dedicated human-approval job) to prompt for human
approval.
2.  If approved, proceed to `publishSummaryTool`.
3.  If rejected, add a message to the `messages` array indicating the
rejection and the reason, and then let the agent loop again, giving the LLM a
chance to regenerate a better summary based on the feedback.

**Hint:**
- You can simulate human approval by using `io.waitforExternalService` with a
custom event that you manually trigger from the dashboard (e.g.,
`human.approval`).
- The payload of this event could include `approved: boolean` and `feedback:
string`.
- Remember to add the human feedback back into the `messages` array for the
LLM to process.

**What to observe/learn:**
- How Trigger.dev's durable execution allows you to pause a workflow for
external human input.
- How the agent can adapt its behavior based on feedback, demonstrating a more
sophisticated agentic loop.

## Common Pitfalls & Troubleshooting for Agents

Building AI agents, especially with complex workflows, can introduce new
challenges.

1.  **Non-Deterministic Behavior:** LLMs are inherently probabilistic. An
agent might take a different path or generate slightly different responses
each time, even with the same input.
    - **Troubleshooting:**
        - **Temperature:** Lower the `temperature` parameter in your LLM calls
(e.g., to 0.0 or 0.1) for more deterministic outputs, especially for critical
decisions.
        - **System Prompt:** Refine your system prompt to be very explicit
about expected behavior, output format, and decision-making logic.
        - **Logging:** Use `io.logger.info` and `io.logger.debug` extensively
to trace the exact messages sent to and received from the LLM, and the tool
calls made. This is invaluable for understanding *why* an agent made a
particular decision.

2.  **Infinite Loops or Max Loop Count Reached:** Agents can sometimes get
stuck in a loop, repeatedly trying the same action or failing to progress.
    - **Troubleshooting:**
        - **`MAX_LOOP_COUNT`:** Always implement a maximum loop count as a
safeguard.
        - **Reflection:** Improve the agent's "reflection" step. Teach the LLM
to recognize when it's stuck or making no progress and to ask for help or
terminate.
        - **Tool Output Clarity:** Ensure tool outputs are clear and concise.
Ambiguous tool outputs can confuse the LLM.
        - **State Management:** Explicitly track the agent's internal state.
If an agent keeps trying the same action, it might not be properly updating
its understanding of the current situation.

3.  **Tool Invocation Failures & Rate Limits:** External APIs (including LLM
APIs) can fail, be slow, or impose rate limits.
```

```
    - **Troubleshooting:**
        - **Trigger.dev Retries:** Leverage Trigger.dev's `retry` options on
`io.runTask` for transient errors.
        - **Error Handling in Tools:** Implement robust `try...catch` blocks
within your tool functions to catch API-specific errors and return meaningful
error messages to the LLM.
        - **Exponential Backoff:** When implementing your own API calls within
tools, use exponential backoff for retries to avoid overwhelming external
services.
        - **Concurrency Limits:** Trigger.dev can manage concurrency for your
jobs, preventing you from hitting external API rate limits too aggressively
from multiple concurrent agent runs.
```

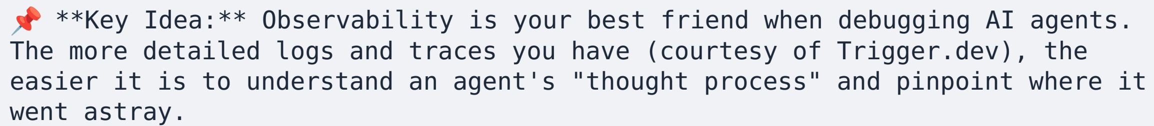

```
📌 **Key Idea:** Observability is your best friend when debugging AI agents.
The more detailed logs and traces you have (courtesy of Trigger.dev), the
easier it is to understand an agent's "thought process" and pinpoint where it
went astray.

## Summary

Congratulations! You've taken a significant step into the world of AI agents
with Trigger.dev.

Here's a recap of what we covered:

- **AI Agents Defined:** We understood agents as systems that Observe, Plan,
Act, and Reflect to achieve goals.
- **Trigger.dev's Role:** We learned how Trigger.dev's durable execution,
retries, and observability provide a robust foundation for building reliable,
long-running agentic workflows.
- **Agentic Job Creation:** You set up a basic Trigger.dev job to orchestrate
LLM calls.
- **Tool Integration:** We enhanced the agent by integrating external tools
using OpenAI's function calling, allowing the agent to interact with the
outside world.
- **Mini-Challenge:** You were tasked with adding a human-in-the-loop approval
step, further demonstrating Trigger.dev's flexibility for complex workflows.
- **Troubleshooting:** We discussed common pitfalls like non-determinism,
infinite loops, and tool failures, along with strategies for addressing them.

By combining the intelligence of LLMs with the reliability of Trigger.dev, you
can build powerful, automated systems that go beyond simple scripts, tackling
complex tasks with resilience and adaptability.

In the next chapter, we'll delve deeper into advanced deployment strategies
and how to effectively monitor your Trigger.dev applications in production.

## References

- [Trigger.dev Documentation](https://trigger.dev/docs)
- [Trigger.dev GitHub Repository](https://github.com/triggerdotdev/
trigger.dev)
- [OpenAI API Documentation](https://platform.openai.com/docs/api-reference)
- [OpenAI Function Calling Guide](https://platform.openai.com/docs/guides/
function-calling)

This page is AI-assisted and reviewed. It references official documentation
and recognized resources where relevant.
```

CHAPTER 08

Human-in-the-Loop & Real-time Updates: Collaborative Workflows

Introduction: The Human Touch in Automated Systems

In the world of AI and automation, achieving fully autonomous systems is often the goal, but not always the best or safest path. Many critical workflows, especially those involving sensitive data, creative output, or high-stakes decisions, benefit immensely from human oversight. This is where **Human-in-the-Loop (HITL)** workflows come into play. They allow automated processes to pause, seek human input, and then continue based on that decision, ensuring accuracy, compliance, and ethical considerations.

Coupled with HITL, **real-time updates** are crucial. When a human is involved, they need immediate feedback on the workflow's status or to quickly provide their input. Real-time updates bridge the gap between durable backend processes and responsive user interfaces, making collaborative workflows feel seamless and efficient.

In this chapter, we'll explore how Trigger.dev empowers you to build robust HITL workflows that can gracefully wait for human interaction and how to design systems that provide real-time updates. You'll learn the core mechanics of pausing and resuming durable jobs based on external events, a foundational skill for creating sophisticated AI agents and collaborative business processes. We assume you're familiar with the basics of creating Trigger.dev jobs and have your development environment set up from previous chapters.

Core Concepts: Weaving Humans into Automated Flows

Building systems that combine the speed and scale of automation with the nuanced judgment of humans requires careful orchestration. Trigger.dev provides the durable execution primitives needed to make this possible without complex state management on your part.

What is Human-in-the-Loop (HITL)?

📌 **Key Idea:** Human-in-the-Loop (HITL) refers to workflows where a human intervenes at specific points to review, approve, or provide input to an otherwise automated process.

Imagine an AI agent generating marketing copy. While the AI can produce content quickly, a human editor needs to review it for tone, accuracy, and brand compliance before it's published. This review step is a classic example of HITL.

Why is HITL so important, especially with AI agents?

- **Accuracy and Quality:** Humans can catch errors, biases, or subtle nuances that AI models might miss.
- **Ethical Oversight:** For sensitive applications, human review ensures decisions align with ethical guidelines and prevent unintended consequences.
- **Compliance and Regulation:** Many industries require human approval for certain actions to meet legal or regulatory standards.
- **Handling Edge Cases:** AI models perform well on typical data, but humans excel at handling novel or ambiguous situations.
- **Continuous Improvement:** Human feedback can be used to retrain and improve AI models over time.

The Role of Real-time Updates

Real-time updates ensure that when a workflow is waiting for human input, the human interface (like a web dashboard or an email notification) reflects the current state accurately and immediately.

For instance, when a job requests human approval, the dashboard should instantly show "Pending Approval." Once the human acts, it should update to "Approved" or "Rejected" without delay. This responsiveness is critical for:

- **User Experience:** No one likes waiting or refreshing pages manually.
- **Transparency:** Users understand where the workflow stands.
- **Efficiency:** Humans can react faster when notified promptly.

Trigger.dev itself is a durable execution platform, meaning it orchestrates backend jobs. While it doesn't directly provide real-time frontend updates (like websockets), it provides the robust eventing mechanism that allows your frontend

to listen for or poll status changes. When a Trigger.dev job sends an event for human interaction, your frontend can be designed to pick up that event or the resulting state change.

Trigger.dev's Approach to HITL: Durable Waits

Trigger.dev makes HITL possible through its durable execution capabilities, specifically the ability for a job to **pause and wait for an external event to resume**. This means:

1. A job can perform some tasks.
2. It can then `await client.waitForEvent()` for a specific event.
3. The workflow is persisted and doesn't consume compute resources while waiting.
4. An external system (like your frontend, an email service, or another backend) can then send an event back to Trigger.dev.
5. When the matching event arrives, the workflow automatically resumes from where it left off, processing the event's payload.

This mechanism ensures that even if your server restarts or goes down, the workflow state is maintained, and it will pick up exactly when the human input arrives.

Here's a simplified flow of a HITL workflow:

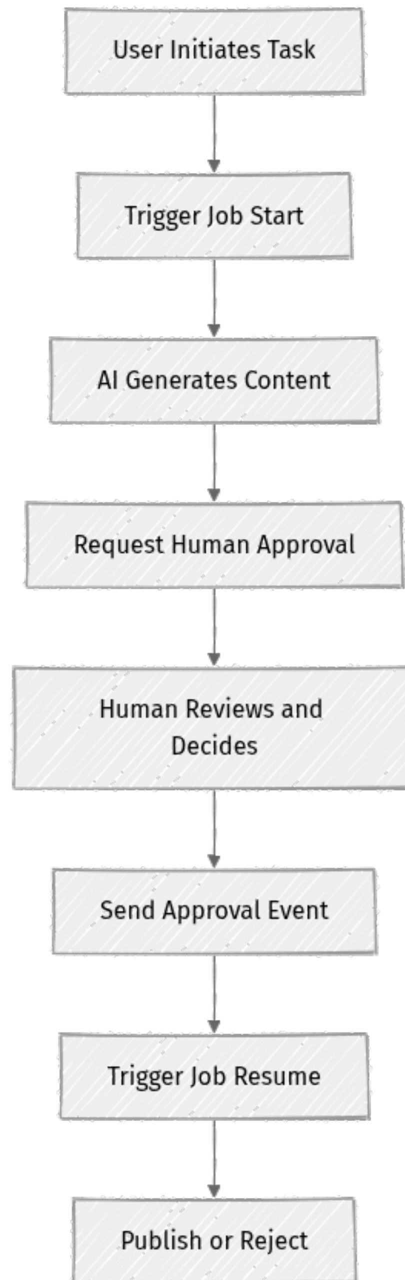

Building a Human Approval Workflow (Step-by-Step)

Let's build a practical example: an AI-generated blog post needs human approval before it's published. We'll simulate the AI generation and the human interaction.

Prerequisites

Make sure you have a Trigger.dev project set up. If you're following along, you should have a `trigger.config.ts` and an `src/jobs` directory.

Step 1: Create a New Job for Content Approval

First, let's create a new job file. We'll call it `src/jobs/contentApproval.ts`.

Create the file:```bash touch src/jobs/contentApproval.ts

````
Open `src/jobs/contentApproval.ts` and add the basic job structure:
```typescript
// src/jobs/contentApproval.ts
import { client } from "../trigger";
import { eventTrigger } from "@trigger.dev/sdk";

client.defineJob({
  id: "content-approval-workflow",
  name: "AI Content Approval Workflow",
  version: "1.0.0",
  trigger: eventTrigger({
    name: "content.generate",
    schema: {
      type: "object",
      properties: {
        topic: { type: "string" },
        draftId: { type: "string" },
      },
      required: ["topic", "draftId"],
      additionalProperties: false,
    },
  }),
  run: async (payload, io, ctx) => {
    // This is where our workflow logic will go
    io.logger.info("Starting AI content approval workflow", { payload });

    // Simulate AI content generation
    await io.wait("waiting-for-ai", 5); // Simulate a 5-second AI generation
    const aiGeneratedContent = `Draft blog post about "${payload.topic}"...
This is placeholder content generated by AI.`;

    io.logger.info("AI content generated.", { content: aiGeneratedContent });

    // We'll add human approval logic here
    // ...
````
````

```
  },
});
```

- **Explanation:**
  - We import `client` and `eventTrigger`.
  - `id` and `name` identify our job.
  - The `trigger` is an `eventTrigger` named `content.generate`. This means our workflow will start when an event with this name is sent to Trigger.dev.
  - The `schema` defines the expected `payload` for this event, which includes `topic` and `draftId`.
  - Inside `run`, we simulate AI content generation using `io.wait`. `io.wait` is a durable way to pause a workflow for a specified duration.

Step 2: Requesting Human Approval

Now, let's introduce the human interaction. The job will "send" an event to notify an external system (which would be your human interface, e.g., a web dashboard or an email system) that approval is needed. Then, it will durably wait for that system to send an approval event back.

Add the following code inside the `run` function, right after `io.logger.info("AI content generated...", { content: aiGeneratedContent });`:

```
    // ... (previous code in run function)

    io.logger.info("Requesting human approval for content.", { draftId: payloa
d.draftId });

    // 🧠 Important: We're sending an event to an *external system*.
    // This event would typically trigger a notification (email, Slack,
dashboard update)
    // in your application, prompting a human to review.
    // The `id` is crucial for linking this specific approval request to the
response.
    const approvalRequestEventId = `content-approval-${payload.draftId}`;

    // You might send this event to a specific queue or topic for your
external system
    // For this example, we're just logging it, but in a real app, this would
    // be `io.sendEvent("external.system.approval.request", { ... })`
    // or an HTTP call to your UI's API.
    io.logger.debug("Simulating sending approval request to external
system.", {
        eventId: approvalRequestEventId,
        content: aiGeneratedContent,
        topic: payload.topic,
    });
```

```
    // ⚡ Quick Note: The `io.waitForEvent` call below is the actual durable
wait.
    // The `io.logger.debug` above is just a placeholder for how you'd notify
    // a human interface.

    // ... (rest of the run function)
```

- **Explanation:**
  - We generate a unique `approvalRequestEventId` using the `draftId`. This ID is vital because it allows us to link the specific approval request to its corresponding response event.
  - We use `io.logger.debug` to simulate sending a notification to an external system. In a real application, this might involve:
    - Calling `io.sendEvent` to a different Trigger.dev job that handles sending emails or Slack messages.
    - Making an `io.runTask` HTTP request to an API endpoint in your frontend application to update a dashboard.
  - The core idea is that something external needs to be notified and given the `approvalRequestEventId` so it can send back the corresponding response.

Step 3: Waiting for Approval

Now for the magic! The job will pause and wait for an event named `content.approved` with a matching `id`.

Add the following code after the `io.logger.debug` call in the `run` function:

```
    // ... (previous code in run function)

    let approvalResult;
    try {
      // 🧠 Important: The job will durably wait here for an event named
"content.approved"
      // with a matching `id` (this links it to our specific approval
request).
      // We also set a timeout of 1 hour (3600 seconds) for the human to
respond.
      approvalResult = await io.waitForEvent<{ approved: boolean; reason?: str
ing }>(
        "human-approval", // This is the unique name for this specific wait
point
        {
          name: "content.approved",
          id: approvalRequestEventId, // Match the ID we generated earlier
          timeoutInSeconds: 3600, // Wait for up to 1 hour
        }
      );
```

```
    } catch (error) {
      if (error instanceof Error && error.message.includes("Timed out")) {
        io.logger.warn("Human approval timed out.", { draftId:
payload.draftId });
        // Handle timeout: e.g., send a reminder, automatically reject, or
escalate
        await io.runTask("send-timeout-alert", async () => {
          // Simulate sending a timeout alert
          io.logger.error(`Approval for draft ${payload.draftId} timed out.`);
        });
        return; // Exit the workflow or take further action
      }
      throw error; // Re-throw other errors
    }

    // ... (rest of the run function)
```

- **Explanation:**
    - `await io.waitForEvent()` is the core function for durable waiting.
    - The first argument, `"human-approval"`, is a unique key for this specific wait point within your workflow. This is useful for observability in the Trigger.dev dashboard.
    - The second argument is an options object:
        - `name: "content.approved"`: This tells Trigger.dev to wait for an incoming event with this specific name.
        - `id: approvalRequestEventId`: This is crucial! It ensures the workflow only resumes when an event with this exact ID is received. This prevents one approval event from accidentally resuming the wrong workflow instance.
        - `timeoutInSeconds: 3600`: If no matching event arrives within 1 hour, `waitForEvent` will throw a timeout error. We wrap this in a `try...catch` block to handle the timeout gracefully.
    - If a timeout occurs, we log a warning, simulate sending an alert, and `return` to stop the workflow or implement further retry/escalation logic.

Step 4: Processing the Decision

Once the `content.approved` event is received (or timed out), the workflow resumes. We can then check the `approvalResult` payload to determine the human's decision.

Add the following code after the `try...catch` block in the `run` function:

```
    // ... (previous code in run function)

    if (approvalResult.payload.approved) {
      io.logger.info("Human approved the content!", { draftId:
payload.draftId });
      // In a real scenario, you'd publish the content here
      await io.runTask("publish-content", async () => {
        io.logger.info(`Publishing content for topic: ${payload.topic}`);
        // Simulate publishing
        await new Promise(resolve => setTimeout(resolve, 2000));
      });
      io.logger.info("Content published successfully.");
    } else {
      io.logger.warn("Human rejected the content.", {
        draftId: payload.draftId,
        reason: approvalResult.payload.reason,
      });
      // In a real scenario, you'd handle rejection (e.g., notify author, send
back for revision)
      await io.runTask("handle-rejection", async () => {
        io.logger.info(`Notifying author about rejection for topic: $
{payload.topic}`);
        // Simulate notification
        await new Promise(resolve => setTimeout(resolve, 1000));
      });
      io.logger.info("Content rejection handled.");
    }

    io.logger.info("AI content approval workflow completed.");
  },
});
```

- **Explanation:**
  - We check `approvalResult.payload.approved`. The `content.approved` event we're waiting for is expected to have a boolean `approved` property in its payload.
  - Based on the `approved` status, we branch the workflow:
    - If `true`, we simulate publishing the content.
    - If `false`, we simulate handling the rejection, potentially including a `reason` from the human.
  - `io.runTask` is used for these side effects, ensuring they are retried if transient failures occur.

The Complete Job Code

For context, here's the full `src/jobs/contentApproval.ts` file:

```
// src/jobs/contentApproval.ts
import { client } from "../trigger";
import { eventTrigger } from "@trigger.dev/sdk";
```

```
client.defineJob({
  id: "content-approval-workflow",
  name: "AI Content Approval Workflow",
  version: "1.0.0",
  trigger: eventTrigger({
    name: "content.generate",
    schema: {
      type: "object",
      properties: {
        topic: { type: "string" },
        draftId: { type: "string" },
      },
      required: ["topic", "draftId"],
      additionalProperties: false,
    },
  }),
  run: async (payload, io, ctx) => {
    io.logger.info("Starting AI content approval workflow", { payload });

    // Simulate AI content generation
    await io.wait("waiting-for-ai", 5); // Simulate a 5-second AI generation
    const aiGeneratedContent = `Draft blog post about "${payload.topic}"...
This is placeholder content generated by AI.`;

    io.logger.info("AI content generated.", { content: aiGeneratedContent });

    io.logger.info("Requesting human approval for content.", { draftId: payloa
d.draftId });

    const approvalRequestEventId = `content-approval-${payload.draftId}`;

    io.logger.debug("Simulating sending approval request to external
system.", {
        eventId: approvalRequestEventId,
        content: aiGeneratedContent,
        topic: payload.topic,
    });

    let approvalResult;
    try {
      approvalResult = await io.waitForEvent<{ approved: boolean; reason?: str
ing }>(
        "human-approval",
        {
          name: "content.approved",
          id: approvalRequestEventId,
          timeoutInSeconds: 3600, // Wait for up to 1 hour
        }
      );
    } catch (error) {
      if (error instanceof Error && error.message.includes("Timed out")) {
        io.logger.warn("Human approval timed out.", { draftId:
payload.draftId });
        await io.runTask("send-timeout-alert", async () => {
          io.logger.error(`Approval for draft ${payload.draftId} timed out.`);
        });
        return;
      }
      throw error;
    }
```

```
    if (approvalResult.payload.approved) {
      io.logger.info("Human approved the content!", { draftId:
payload.draftId });
      await io.runTask("publish-content", async () => {
        io.logger.info(`Publishing content for topic: ${payload.topic}`);
        await new Promise(resolve => setTimeout(resolve, 2000));
      });
      io.logger.info("Content published successfully.");
    } else {
      io.logger.warn("Human rejected the content.", {
        draftId: payload.draftId,
        reason: approvalResult.payload.reason,
      });
      await io.runTask("handle-rejection", async () => {
        io.logger.info(`Notifying author about rejection for topic: $
{payload.topic}`);
        await new Promise(resolve => setTimeout(resolve, 1000));
      });
      io.logger.info("Content rejection handled.");
    }

    io.logger.info("AI content approval workflow completed.");
  },
});
```

Step 5: Simulating Real-time Feedback and Interaction

To test this, you need to:

1. **Start your Trigger.dev development server:**

```
npx trigger.dev@v4-beta dev
```

1. **Trigger the initial event:** Go to the Trigger.dev dashboard (usually `<http://localhost:8888`>), find your `AI Content Approval Workflow` job, and click "Run Job". Provide a payload like this:

```
{
  "topic": "The Future of AI in Education",
  "draftId": "blog-post-123"
}
```

```
The job will start, simulate AI generation, and then enter a "WAITING" state.
You'll see this in the Trigger.dev dashboard.
```

1. **Simulate human approval:** While the job is waiting, you need to send an event named `content.approved` with the correct `id` back to Trigger.dev.
    - The `id` should be `content-approval-blog-post-123` (based on our `draftId`).
    - The payload needs `{ "approved": true }` or `{ "approved": false, "reason": "Needs more examples." }`.

    You can send this event using the Trigger.dev dashboard's "Send Event" feature, or programmatically. Here's how you'd send an event programmatically (e.g., from an API endpoint in your Next.js app that a human UI button calls):

```
    // Example of how an external system (e.g., a Next.js API route) would
send the event
    import { TriggerClient } from "@trigger.dev/sdk";

    // You would initialize this client with your API key and project ID
    // For local dev, ensure TRIGGER_API_KEY and TRIGGER_PROJECT_ID are set in
your .env
    const triggerClient = new TriggerClient({
      id: "my-app-worker", // A unique ID for your client sending events
      apiKey: process.env.TRIGGER_API_KEY!,
    });

    async function sendApprovalEvent(draftId: string, approved: boolean, reaso
n?: string) {
      const eventId = `content-approval-${draftId}`;
      await triggerClient.sendEvent({
        name: "content.approved",
        id: eventId, // CRITICAL: This links back to the waiting workflow
        payload: {
          approved: approved,
          reason: reason,
        },
      });
      console.log(`Sent approval event for draft ${draftId}. Approved: ${appro
ved}`);
    }

    // Call this function when a human approves:
    // sendApprovalEvent("blog-post-123", true);

    // Call this function when a human rejects:
    // sendApprovalEvent("blog-post-123", false, "Needs more real-world
examples.");
```

```
- **Explanation:**
    - We initialize a `TriggerClient` with your project's `apiKey` and `id`.
    - `triggerClient.sendEvent` is used to send the `content.approved` event.
    - The `id` field is set to `content-approval-blog-post-123`, which matches
the `id` our workflow is waiting for. This is how Trigger.dev knows which
specific workflow instance to resume.
    - The `payload` contains the human's decision (`approved: boolean`).
```

After sending the approval event, go back to the Trigger.dev dashboard and observe your workflow. It should immediately resume from its "WAITING" state and complete based on the approval decision you sent. This immediate resumption demonstrates the "real-time" aspect of the workflow reacting to external input.

Mini-Challenge: Add Rejection Handling

Your turn!

Challenge: Modify the `contentApproval.ts` workflow to include an additional step when the content is rejected. Specifically, if the content is rejected, instead of just logging a message, durable wait for a "revision completed" event, then restart the AI generation and approval process. This simulates a cycle where a human sends content back for revision.

Hint:

- Inside the `else` block (where content is rejected), you'll need another `io.waitForEvent` for a `content.revision.completed` event, again using the `approvalRequestEventId` to link it.
- After the revision event is received, you can re-run the AI generation and re-request approval. You might need to wrap the AI generation and approval request logic in a function to easily call it again.
- Consider adding a `revisionCount` to the job payload to prevent infinite loops.

What to observe/learn: You'll see how to create complex, multi-stage human-in-the-loop workflows that can loop back to earlier stages based on human decisions. You'll also further solidify your understanding of `io.waitForEvent` and event IDs.

Common Pitfalls & Troubleshooting

Working with human-in-the-loop and real-time updates introduces a few common challenges:

1. **Timeout Management:**
    - **Pitfall:** Not setting a `timeoutInSeconds` for `io.waitForEvent`, or setting it too short/long. If there's no timeout, a workflow could wait forever if the human never responds or the external event is never sent. If it's too short, it might time out prematurely.
    - **Troubleshooting:** Always consider the expected human response time. Implement robust error handling for `TimeoutError` as shown in our example. You might want to send reminders before the timeout, escalate to another human, or automatically take a default action. Monitor your Trigger.dev dashboard for workflows stuck in "WAITING" states.
2. **Event ID Mismatches:**
    - **Pitfall:** The `id` passed to `io.waitForEvent` does not exactly match the `id` of the incoming event. This is a very common mistake. The workflow will never resume because it's waiting for a specific identifier that never arrives.
    - **Troubleshooting:** Double-check that the unique `id` generated and used in `io.waitForEvent` is exactly the same as the `id` sent with `triggerClient.sendEvent`. Use descriptive IDs that clearly link to the specific workflow instance and context (e.g., `content-approval-${draftId}`).
3. **External System Reliability:**
    - **Pitfall:** The external system responsible for sending the approval event (e.g., your frontend API, an email service) fails or doesn't send the event correctly.
    - **Troubleshooting:** Ensure your external system has proper error handling and logging when it attempts to send events to Trigger.dev. Use `io.runTask` for any external calls from your Trigger.dev job (e.g., sending an email notification) so that these actions are retried if they fail. Implement observability in your external system to confirm events are being sent.

4. **Debugging Durable Waits:**
    - **Pitfall:** It can be confusing to debug a workflow that's paused.
    - **Troubleshooting:** The Trigger.dev dashboard is your best friend. When a job is in a "WAITING" state, you can see which `io.waitForEvent` call it's on, what event `name` and `id` it's waiting for, and how much time is left on its timeout. This provides crucial context for debugging.

Summary

In this chapter, we've explored the essential concepts and practical implementation of Human-in-the-Loop (HITL) workflows and how to enable real-time updates using Trigger.dev.

Here are the key takeaways:

- **HITL Importance:** Human intervention is critical for accuracy, ethics, compliance, and handling edge cases in many AI and automated systems.
- **Durable Waiting:** Trigger.dev's `io.waitForEvent()` function allows workflows to durably pause, persist their state, and resume only when a specific external event arrives.
- **Event IDs:** Unique `id` values are crucial for linking a specific `io.waitForEvent()` call to its corresponding incoming event, preventing cross-talk between workflow instances.
- **Timeout Handling:** Implementing `timeoutInSeconds` and handling `TimeoutError` in `io.waitForEvent()` is vital for robust workflows that don't get stuck indefinitely.
- **Real-time Interaction:** While Trigger.dev orchestrates the backend, its eventing model allows your external systems (like web UIs) to send and receive events, enabling responsive, real-time feedback for human users.

You now have the tools to design and implement collaborative workflows where AI and humans work together seamlessly. In the next chapter, we'll dive into integrating Trigger.dev with other parts of your modern application stack, specifically focusing on Next.js, and explore how to deploy and manage these powerful workflows in production.

References

- Trigger.dev Documentation: https://trigger.dev
- Trigger.dev GitHub Repository: https://github.com/triggerdotdev/trigger.dev
- Trigger.dev `io.waitForEvent` documentation (search for `waitForEvent`): https://trigger.dev/docs/sdk/reference/io

---

This page is AI-assisted and reviewed. It references official documentation and recognized resources where relevant.

CHAPTER 09

Advanced Integrations: Understanding MCP & Custom Connectors

Introduction

As you build increasingly sophisticated AI agents and automated workflows, you'll inevitably encounter the need to connect to a wider array of services than any platform can offer out-of-the-box. This is where advanced integrations become crucial. You might need to interact with a niche third-party API, a legacy internal system, or perhaps a highly specialized AI model hosted in a unique environment.

This chapter dives into how Trigger.dev empowers you to go beyond its standard integrations. We'll explore the concept of the Managed Connector Platform (MCP) and, more importantly, guide you through building your own custom connectors. Mastering this skill allows your Trigger.dev workflows to truly become the central nervous system for all your operations, regardless of how obscure or proprietary your external services might be.

To get the most out of this chapter, you should be comfortable with basic Trigger.dev workflow creation, understanding of event-driven architectures, and have a foundational grasp of API interactions and TypeScript. We're building on the knowledge gained in previous chapters, so prepare to expand your horizons!

Core Concepts: Extending Trigger.dev's Reach

Trigger.dev provides a robust foundation for orchestrating workflows and AI agents. But its true power lies in its ability to seamlessly integrate with virtually any external service. This is achieved through its connector ecosystem, which includes both a Managed Connector Platform and the flexibility to create custom connectors.

The Managed Connector Platform (MCP)

What if you could connect to dozens of popular SaaS applications without worrying about API keys, authentication flows, rate limits, or error handling? That's precisely the problem the Managed Connector Platform (MCP) solves.

What it is: The Trigger.dev Managed Connector Platform (MCP) is a system that provides and manages a collection of pre-built, production-ready integrations with popular third-party services. Think of it as a curated library of "smart adapters" for common APIs like Stripe, GitHub, Slack, OpenAI, and more. These are often developed and maintained by the Trigger.dev team or trusted partners.

Why it exists: Integrating with external APIs is often tedious. Each service has its own authentication method (API keys, OAuth, etc.), rate limiting, error codes, and data structures. The MCP abstracts away this complexity, offering a standardized, simplified interface within your Trigger.dev workflows. This significantly reduces development time and the potential for integration-specific bugs.

How it works: When you use an MCP-provided connector (e.g., `github.runsTask`), Trigger.dev handles the underlying HTTP requests, authentication (using secrets you provide), retries for transient errors, and often transforms the data into a consistent format. You simply call a function, and the MCP ensures it reaches the external service reliably.

⚡ **Real-world insight:** In a production system, relying on MCP connectors for common services means your team can focus on core business logic rather than building and maintaining dozens of API clients. It's a huge time-saver and reliability booster.

The Need for Custom Connectors

While the MCP covers many popular services, the world of software is vast. You'll inevitably encounter scenarios where a pre-built connector doesn't exist or doesn't precisely fit your needs.

When built-in isn't enough:

- **Proprietary Internal Systems:** Your company might have its own APIs for CRM, ERP, or custom data stores that are not publicly exposed.
- **Niche Third-Party Services:** You might be working with a highly specialized vendor API that isn't widely adopted enough for an MCP connector.
- **Unique Data Transformations:** You might need to perform specific data manipulation before or after interacting with an API, which a generic connector might not support.
- **Custom Authentication Flows:** Some services require complex multi-step authentication that differs from standard API key or OAuth patterns.

Benefits of Custom Connectors:

- **Tailored Logic:** You have full control over the interaction, allowing for highly specific requests and responses.
- **Seamless Integration:** Your proprietary systems can become first-class citizens in your Trigger.dev workflows.
- **Enhanced Security:** You control how secrets are used and transmitted, aligning with internal security policies.
- **Future-Proofing:** You can adapt your connector as the external API evolves, without waiting for a platform update.

Anatomy of a Custom Trigger.dev Connector

A custom Trigger.dev connector is essentially a TypeScript module that defines how your workflows can interact with an external service. It typically consists of:

- **ID and Name:** Unique identifiers for your connector.
- **Schemas:** Define the configuration options and authentication secrets needed for the connector.
- **Actions:** Functions that perform specific operations (e.g., `createUser`, `sendNotification`, `fetchData`). These are the methods your workflows will call.
- **Triggers (Optional):** Mechanisms for the external service to initiate a Trigger.dev workflow (e.g., via webhooks).

Let's visualize the core components of a custom connector:

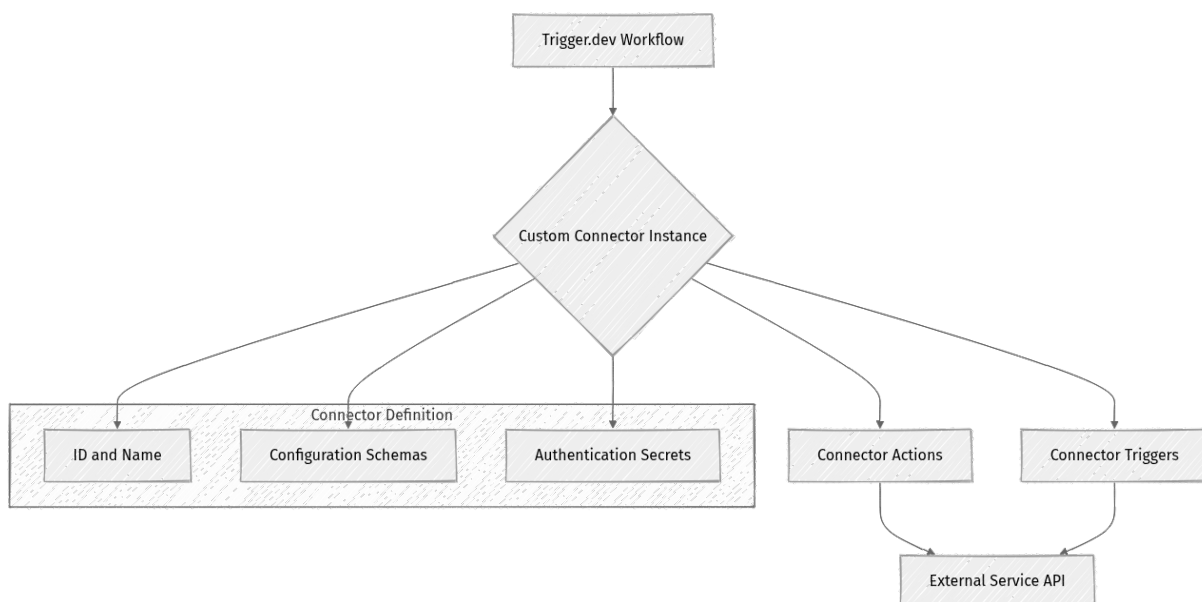

- **Trigger.dev Workflow:** Your actual workflow code that orchestrates tasks.

- **Custom Connector Instance:** How your workflow interacts with your custom connector.
- **Connector Actions:** Functions within your connector that perform operations on the external service.
- **Connector Triggers:** How events from the external service can start your Trigger.dev workflows.
- **External Service API:** The actual API endpoint of the service you're integrating with.
- **Connector Definition:** The metadata and configuration details that describe your connector.

Building a Custom Connector: A Step-by-Step Guide

For this guide, we'll create a simple custom connector for a hypothetical "Simple Chat Service" that allows us to send messages.

Step 1: Project Setup

We'll assume you already have a Trigger.dev project initialized. If not, make sure to use `v4-beta`. As of 2026-05-20, v3 is the current stable, but v4 is expected to go GA around May/June 2026.

```
npx trigger.dev@v4-beta init
```

Inside your Trigger.dev project, navigate to your `src` directory. We'll create a new folder for our custom connectors.

```
# In your project root
mkdir -p src/connectors/simple-chat
touch src/connectors/simple-chat/index.ts
```

This creates a dedicated directory for our `simple-chat` connector and an `index.ts` file where its logic will reside.

Step 2: Defining the Connector Interface

First, we need to define the basic structure of our connector using Trigger.dev's `defineConnector` helper. This includes its `id`, `name`, and an optional `schemas` object for configuration.

Open `src/connectors/simple-chat/index.ts` and add the following:

```
// src/connectors/simple-chat/index.ts
import { defineConnector } from "@trigger.dev/sdk";
import { z } from "zod"; // We'll use Zod for schema validation

// Define the configuration schema for our connector
// This allows users to provide an API key when setting up the connector
const SimpleChatSchema = z.object({
  apiKey: z.string(), // A string for our API key
  // We could add more config options here, like a default channel ID
});

export const simpleChat = defineConnector({
  id: "simple-chat", // Unique ID for the connector
  name: "Simple Chat Service", // Human-readable name
  version: "1.0.0", // Version of your connector
  schemas: {
    // This schema will be used to validate the config object when the
connector is initialized
    // It also helps Trigger.dev understand what secrets/config the connector
needs
    config: SimpleChatSchema,
  },
});
```

Explanation:

- `defineConnector`: This is the core function from `@trigger.dev/sdk` that helps you build a connector.
- `zod`: A TypeScript-first schema declaration and validation library. Trigger.dev uses Zod internally for schema validation, making it a natural fit for defining connector configurations.
- `SimpleChatSchema`: Defines that our connector needs an `apiKey` which must be a string. This `apiKey` will be provided by the user when they connect to the `Simple Chat Service` in their Trigger.dev dashboard or directly in code.
- `id`: A unique identifier for your connector (e.g., `simple-chat`).
- `name`: A user-friendly name that will appear in the Trigger.dev UI.
- `version`: A semantic version for your connector, useful for managing updates.
- `schemas.config`: This tells Trigger.dev what configuration properties your connector expects.

Step 3: Implementing Actions

Actions are the core functionality of your connector. They define what operations your workflows can perform on the external service. Let's add an action to send a message.

Update `src/connectors/simple-chat/index.ts`:

```
// src/connectors/simple-chat/index.ts
import { defineConnector } from "@trigger.dev/sdk";
import { z } from "zod";

const SimpleChatSchema = z.object({
  apiKey: z.string(),
});

export const simpleChat = defineConnector({
  id: "simple-chat",
  name: "Simple Chat Service",
  version: "1.0.0",
  schemas: {
    config: SimpleChatSchema,
  },
}).defineAction({ // We're adding an action here!
  id: "sendMessage", // Unique ID for this action
  name: "Send Message", // Human-readable name for the action
  input: z.object({ // Schema for the input parameters of this action
    channel: z.string(),
    message: z.string(),
  }),
  output: z.object({ // Schema for the output (return value) of this action
    success: z.boolean(),
    messageId: z.string().optional(),
    error: z.string().optional(),
  }),
  run: async (payload, { config, ctx }) => {
    // This is the actual function that runs when the action is called
    // `payload` contains the input defined in `input` schema
    // `config` contains the connector's configuration (our apiKey)
    // `ctx` provides context like logger

    ctx.logger.info(`Sending message to channel ${payload.channel}`);

    try {
      // In a real scenario, you'd make an HTTP request to your Simple Chat
Service here
      // For this example, we'll simulate an API call
      const response = await fetch("https://api.simplechat.com/v1/message", {
        method: "POST",
        headers: {
          "Content-Type": "application/json",
          "Authorization": `Bearer ${config.apiKey}`, // Use the API key from
the config
        },
        body: JSON.stringify({
          channel: payload.channel,
          text: payload.message,
        }),
      });

      if (!response.ok) {
        const errorData = await response.json();
        throw new Error(`Failed to send message: ${response.status} - ${errorD
ata.message}`);
      }
```

```
      const data = await response.json();
      return { success: true, messageId: data.id };

    } catch (error) {
      ctx.logger.error("Error sending message:", error);
      return { success: false, error: (error as Error).message };
    }
  },
});
```

Explanation:

- `.defineAction()`: This method is chained to `defineConnector` to add an action.
- `id` and `name`: Unique identifier and human-readable name for this specific action.
- `input`: A Zod schema defining the expected arguments for this action (e.g., `channel` and `message`). Trigger.dev will validate these inputs before executing `run`.
- `output`: A Zod schema defining the expected return value of this action. This helps with type safety and understanding what data you'll get back.
- `run: async (payload, { config, ctx }) => { ... }`: This is the asynchronous function that executes the action.
  - `payload`: An object containing the validated input parameters (`channel`, `message`).
  - `config`: An object containing the connector's configuration, which includes our `apiKey`.
  - `ctx`: Provides useful utilities like a `logger` for debugging.
- **Simulated API Call:** We've included a `fetch` call to `<https://api.simplechat.com/v1/message`>. In a real application, you would replace this with the actual API endpoint and client logic for your external service. Notice how `config.apiKey` is used for authentication.
- Error Handling: The `try...catch` block ensures that network errors or API errors are caught and returned gracefully.

Step 4: Integrating the Custom Connector into a Workflow

Now that our `simpleChat` connector is defined, we can use its `sendMessage` action in any Trigger.dev workflow.

First, make sure to install `zod` if you haven't already:```bash npm install zod

or

yarn add zod

```
Next, open one of your workflow files (e.g., `src/jobs/my-first-workflow.ts`)
and modify it to use the new connector.
```typescript
// src/jobs/my-first-workflow.ts
import { client, eventTrigger } from "../trigger"; // Ensure eventTrigger is
imported
import { simpleChat } from "../connectors/simple-chat"; // Import our custom
connector
import { z } from "zod"; // Import z for schema definition

// Initialize the custom connector instance
// You would typically store your API key securely, e.g., in environment
variables
// For local development, you can put it in a .env file or directly here for
testing.
// In production, use Trigger.dev's secret management.
const simpleChatService = simpleChat.init({
  apiKey: process.env.SIMPLE_CHAT_API_KEY || "your_dev_api_key_here",
});

client.defineJob({
  id: "send-chat-message",
  name: "Send Chat Message to Custom Service",
  version: "1.0.0",
  integrations: {
    simpleChatService, // Register our custom connector instance
  },
  trigger: eventTrigger({ // Using eventTrigger for a straightforward manual
test
    name: "manual.send.message",
    schema: z.object({
      channel: z.string(),
      message: z.string(), // The payload will have a 'message' field
    }),
  }),
  run: async (payload, io, ctx) => {
    // Use the custom connector's action
    const result = await io.simpleChatService.sendMessage(
      "send-message-action", // Unique key for this step
      {
        channel: payload.channel,
        message: payload.message, // Use payload.message from the trigger
      }
    );

    if (result.success) {
      io.logger.info(`Message sent successfully! Message ID: $
{result.messageId}`);
    } else {
      io.logger.error(`Failed to send message: ${result.error}`);
      throw new Error(`Chat service error: ${result.error}`);
    }

    return { status: "message_processed", messageId: result.messageId };
```
```

```
  },
});
```

Explanation:

- `import { client, eventTrigger } from "../trigger";`: We import `eventTrigger` alongside `client` from our `trigger.ts` file.
- `import { simpleChat } from "../connectors/simple-chat";`: We import our defined connector.
- `simpleChat.init({ apiKey: ... })`: We initialize an instance of our connector, providing the necessary configuration (our `apiKey`). In a real application, `process.env.SIMPLE_CHAT_API_KEY` would fetch the key from your environment variables, which is the recommended way to handle secrets.
- `integrations: { simpleChatService, }`: We register our initialized connector instance with the job. This makes it available via `io.simpleChatService`.
- `trigger: eventTrigger(...)`: We use `eventTrigger` to define a simple, manually callable trigger. The `schema` specifies that the event payload should contain `channel` and `message` properties.
- `io.simpleChatService.sendMessage(...)`: Inside the `run` function, we call our custom action. The first argument is a unique key for this step (important for observability), and the second is the input payload matching our `sendMessage` action's schema. Notice how `payload.message` is consistently used here, matching the `eventTrigger` schema.

Step 5: Implementing Triggers (Conceptual)

While actions allow your workflows to call external services, triggers allow external services to start your workflows. Implementing a trigger for a custom connector involves defining how Trigger.dev listens for events from your external service. This is often done via webhooks.

In your `src/connectors/simple-chat/index.ts`, you could extend `defineConnector` with a `defineTrigger` block.

```
// ... existing code ...

export const simpleChat = defineConnector({
  id: "simple-chat",
  name: "Simple Chat Service",
  version: "1.0.0",
  schemas: {
    config: SimpleChatSchema,
```

```
  },
})
.defineAction({
  // ... sendMessage action ...
})
.defineTrigger({ // Adding a trigger!
  id: "message.received", // Unique ID for this trigger
  name: "Message Received", // Human-readable name
  input: z.object({ // Schema for the payload the external service sends
    channel: z.string(),
    sender: z.string(),
    text: z.string(),
    timestamp: z.string().datetime(),
  }),
  // For a real webhook trigger, you'd configure how Trigger.dev receives and
validates it.
  // This might involve a special endpoint URL provided by Trigger.dev for
your connector,
  // and your external service would send POST requests to it.
  // The 'run' function here would typically just return the validated
payload.
  // The actual webhook endpoint creation and listening is handled by
Trigger.dev internally
  // when you define a trigger like this.
  run: async (payload, { ctx }) => {
    ctx.logger.info(`Received new message from Simple Chat Service: $
{payload.text}`);
    return payload; // Return the payload to be used by the workflow
  },
});
```

Explanation:

- `.defineTrigger()`: Similar to `defineAction`, but for incoming events.
- `id` and `name`: Identifiers for this specific trigger.
- `input`: The Zod schema for the expected payload from the external service.
- `run`: For triggers, the `run` function typically just validates and returns the incoming payload, allowing the workflow to process it. The magic of listening for HTTP requests (webhooks) is handled by Trigger.dev when you deploy this connector.

⚡ **Quick Note:** Setting up a real webhook for a custom service often requires configuring that service to send POST requests to a specific URL provided by Trigger.dev. This URL is generated when your Trigger.dev project is deployed and your connector is registered.

Mini-Challenge: Extend Your Custom Connector

Now it's your turn to expand the functionality of our `simpleChat` connector!

Challenge: Add a new action to the `simpleChat` connector called `getChannelInfo`. This action should take a `channelId` (string) as input and return mock information about that channel, such as `name` (string), `memberCount` (number), and `topic` (string). Simulate an API call with a `setTimeout` to mimic network latency.

Hint:

1. Add another `.defineAction()` block after `sendMessage`.
2. Define `id`, `name`, `input` (for `channelId`), and `output` (for channel info).
3. Inside the `run` function, use `setTimeout` to delay the response and return a hardcoded object matching your `output` schema.

What to observe/learn: You'll practice defining inputs and outputs for actions, simulating asynchronous operations, and extending the capabilities of your custom connector.

```
// Add this inside src/connectors/simple-chat/index.ts,
// after the sendMessage action and before any other code.
.defineAction({
  id: "getChannelInfo",
  name: "Get Channel Information",
  input: z.object({
    channelId: z.string(),
  }),
  output: z.object({
    name: z.string(),
    memberCount: z.number(),
    topic: z.string().optional(),
    found: z.boolean(),
  }),
  run: async (payload, { config, ctx }) => {
    ctx.logger.info(`Fetching info for channel ID: ${payload.channelId}`);

    // Simulate an API call with latency
    await new Promise(resolve => setTimeout(resolve, 500)); // 500ms delay

    if (payload.channelId === "general") {
      return {
        name: "General Chat",
        memberCount: 123,
        topic: "Discussions about everything!",
        found: true,
      };
    } else if (payload.channelId === "dev-team") {
      return {
        name: "Development Team",
        memberCount: 15,
        topic: "Daily standups and coding questions.",
        found: true,
      };
    } else {
      return {
        name: `Channel ${payload.channelId}`,
```

```
        memberCount: 0,
        topic: undefined,
        found: false,
      };
    }
  },
});
```

Once you've added this action, try calling it from a new job or within your existing `send-chat-message` job to verify it works!

```
// Example of using the new action in a job
// src/jobs/get-channel-info-job.ts (create a new file for this job)
import { client, eventTrigger } from "../trigger";
import { simpleChat } from "../connectors/simple-chat";
import { z } from "zod";

// Initialize the custom connector instance
const simpleChatService = simpleChat.init({
  apiKey: process.env.SIMPLE_CHAT_API_KEY || "your_dev_api_key_here",
});

client.defineJob({
  id: "get-chat-channel-info",
  name: "Get Custom Chat Channel Info",
  version: "1.0.0",
  integrations: {
    simpleChatService,
  },
  trigger: eventTrigger({
    name: "manual.get.channel.info",
    schema: z.object({
      channelId: z.string(),
    }),
  }),
  run: async (payload, io, ctx) => {
    const channelInfo = await io.simpleChatService.getChannelInfo(
      "get-channel-info-step",
      { channelId: payload.channelId }
    );

    if (channelInfo.found) {
      io.logger.info(`Channel Name: ${channelInfo.name}`);
      io.logger.info(`Members: ${channelInfo.memberCount}`);
      if (channelInfo.topic) {
        io.logger.info(`Topic: ${channelInfo.topic}`);
      }
    } else {
      io.logger.warn(`Channel with ID ${payload.channelId} not found.`);
    }

    return channelInfo;
  },
});
```

Common Pitfalls & Troubleshooting

Building custom integrations can sometimes be tricky. Here are a few common issues and how to approach them:

- **Authentication Failures:**
  - **Problem:** Your custom connector fails with 401 Unauthorized or 403 Forbidden errors when interacting with the external service.
  - **Debugging:**
    - Double-check your `apiKey` (or other credentials) in your `.env` file or Trigger.dev secrets. Is it correct and does it have the necessary permissions?
    - Ensure your `Authorization` header format is correct (e.g., `Bearer YOUR_KEY`, `Basic BASE64_ENCODED_CREDENTIALS`).
    - Verify the API key hasn't expired or been revoked on the external service's dashboard.
    - ⚡ **Real-world insight:** Always use Trigger.dev's secret management for production API keys, not hardcoded values or `.env` files that might not be deployed securely.
- **Schema Mismatches (Input/Output Validation):**
  - **Problem:** Your action fails because the input `payload` doesn't match your `input` Zod schema, or your `run` function returns data that doesn't match your `output` Zod schema.
  - **Debugging:**
    - Carefully compare the data you're passing to the action with the `input` schema you defined. The error message from Zod will usually pinpoint the exact field causing the issue.
    - Ensure the data returned by your `run` function (the actual `return { ... }` object) perfectly matches your `output` schema, including handling optional fields (`.optional()`).
    - Trigger.dev's runtime will provide helpful error messages indicating where schema validation failed, often with a clear path to fix.

- **Rate Limiting from External APIs:**
  - **Problem:** Your connector works for a few calls, then starts failing with 429 Too Many Requests errors from the external API.
  - **Debugging:**

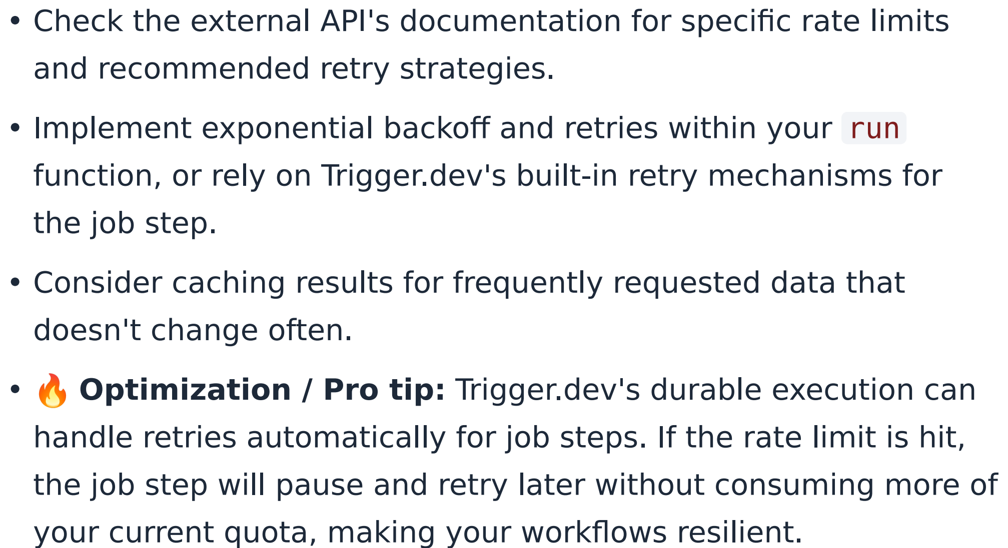

    - Check the external API's documentation for specific rate limits and recommended retry strategies.
    - Implement exponential backoff and retries within your `run` function, or rely on Trigger.dev's built-in retry mechanisms for the job step.
    - Consider caching results for frequently requested data that doesn't change often.
    - 🔥 **Optimization / Pro tip:** Trigger.dev's durable execution can handle retries automatically for job steps. If the rate limit is hit, the job step will pause and retry later without consuming more of your current quota, making your workflows resilient.
- **Debugging Connector Logic:**
  - **Problem:** Your `run` function isn't behaving as expected, but there are no obvious errors.
  - **Debugging:**
    - Use `ctx.logger.info()` and `ctx.logger.error()` extensively within your `run` function. These logs will appear in the Trigger.dev dashboard for the specific job run, providing crucial insights into execution flow and variable states.
    - During local development, you can use `console.log()` to see output directly in your terminal.
    - Break down complex logic into smaller, testable functions to isolate potential issues.

Summary

In this chapter, we've taken a significant leap forward in understanding the power and flexibility of Trigger.dev's integration capabilities.

Here are the key takeaways:

- **MCP (Managed Connector Platform):** Provides ready-to-use integrations for popular services, abstracting away API complexity and streamlining development.
- **Custom Connectors:** Essential for integrating with proprietary, niche, or highly customized external services, offering complete control over API interactions.
- **Connector Anatomy:** Custom connectors are defined by an `id`, `name`, `schemas` for configuration, and can include `actions` (for calling the external service) and `triggers` (for the external service to initiate workflows).
- **Step-by-Step Implementation:** We learned how to define a connector using `defineConnector`, add actions with `defineAction`, specify `input` and `output` schemas, and implement the `run` function for actual API interaction.
- **Integration:** Custom connectors are initialized and registered with a Trigger.dev job, making their actions available via the `io` object.
- **Troubleshooting:** Common issues like authentication failures, schema mismatches, and rate limiting can be resolved through careful configuration, robust logging, and leveraging Trigger.dev's retry mechanisms.

You now have the knowledge to extend Trigger.dev's reach to virtually any API, making your workflows incredibly powerful and adaptable. In the next chapter, we'll delve into deploying your Trigger.dev projects to production environments, ensuring your robust workflows and custom integrations are ready for real-world traffic.

References

- Trigger.dev Documentation: https://trigger.dev/docs
- Trigger.dev GitHub Repository: https://github.com/triggerdotdev/trigger.dev
- Zod GitHub Repository: https://github.com/colinhacks/zod

This page is AI-assisted and reviewed. It references official documentation and recognized resources where relevant.

CHAPTER 10

Self-Hosting Trigger.dev: Taking Full Control (Advanced)

Imagine needing ultimate control over your workflow execution engine. Perhaps strict data residency, specific security policies, or a desire for deep infrastructure customization dictates your approach. While Trigger.dev offers a robust managed cloud service, for advanced users and specific enterprise scenarios, self-hosting becomes a powerful, indispensable option.

This chapter dives into the complex yet rewarding world of self-hosting Trigger.dev. We'll dissect its underlying architecture, guide you through a local setup using Docker Compose, and discuss critical considerations for deploying it securely and scalably in a production environment. Be prepared for a hands-on journey that gives you complete command over your workflow infrastructure.

This guide assumes you have a solid understanding of Docker, container orchestration concepts, and database management. We'll focus on Trigger.dev `v4-beta`, which is slated for General Availability around May/June 2026. While `v3` is the current stable release, `v4-beta` offers the latest features and architectural patterns for self-hosting. We'll emphasize using the `v4-beta` branch for the most current experience.

⚡ **Quick Note:** The prompt mentioned "MCP integration." Based on available documentation (checked 2026-05-20), "MCP" doesn't refer to a standard, self-hostable component of Trigger.dev itself. If it refers to a "Master Control Program" for AI agents, this is typically an architectural pattern or system you would build using Trigger.dev's robust workflow capabilities, rather than a separate Trigger.dev service to self-host.

Understanding the Self-Hosting Architecture

Self-hosting Trigger.dev means taking full responsibility for all the underlying services that power its durable execution and AI agent capabilities. It's not a single application; it's a collection of interconnected, resilient, and scalable services. Understanding these components is the first step toward effective management.

Core Components of Trigger.dev

Trigger.dev orchestrates jobs and workflows across several key services. When you self-host, you deploy and manage each of these essential parts:

1. **Trigger.dev API:**
    - **What it is:** The primary entry point for your client applications (your `trigger.dev` project).
    - **Why it's important:** It receives workflow definitions, ingests events, schedules jobs, and acts as the central hub for all client-side interactions.
    - **How it functions:** It exposes HTTP endpoints for your client projects to communicate with, validating requests and pushing events into the queuing system.
2. **Trigger.dev Worker:**
    - **What it is:** The execution engine for your actual workflow code.
    - **Why it's important:** Workers pick up jobs from the queue, execute the defined steps (including calls to external services), handle retries, and report the status back to the API.
    - **How it functions:** It continuously polls the job queue, fetches tasks, and runs your TypeScript/JavaScript workflow functions in a sandboxed environment. You'll typically run multiple workers for fault tolerance and scalability.
3. **Trigger.dev Dashboard:**
    - **What it is:** The user interface for monitoring and managing your Trigger.dev projects and workflow runs.
    - **Why it's important:** Provides observability into your system, allowing you to view runs, inspect logs, manage API keys, and debug issues.
    - **How it functions:** It's a web application that interacts with the Trigger.dev API and PostgreSQL database to display real-time and historical data.

4. **PostgreSQL Database:**
    - **What it is:** The central persistent storage for Trigger.dev.
    - **Why it's important:** It's critical for durable execution, storing all essential state including workflow definitions, job statuses, run histories, event logs, and user data. This ensures workflows can resume after interruptions.
    - **How it functions:** All Trigger.dev services interact with PostgreSQL to read and write state, ensuring consistency and reliability across distributed operations.
5. **Redis Cache/Queue:**
    - **What it is:** A high-performance in-memory data store used as a message broker and cache.
    - **Why it's important:** It facilitates real-time communication between the API and workers, manages job queues (ensuring jobs are picked up by available workers), and stores transient data for quick access.
    - **How it functions:** The API enqueues jobs into Redis, and workers dequeue them. Redis's speed is crucial for minimizing latency in job processing.

These components work in concert to provide the resilient and scalable workflow execution that Trigger.dev is known for.

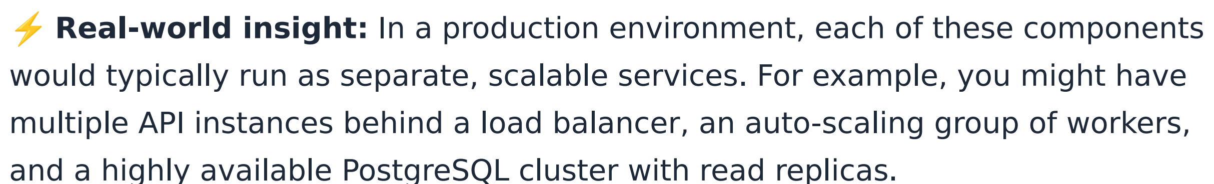

⚡ **Real-world insight:** In a production environment, each of these components would typically run as separate, scalable services. For example, you might have multiple API instances behind a load balancer, an auto-scaling group of workers, and a highly available PostgreSQL cluster with read replicas.

Here's a simplified view of how these components interact in a self-hosted setup:

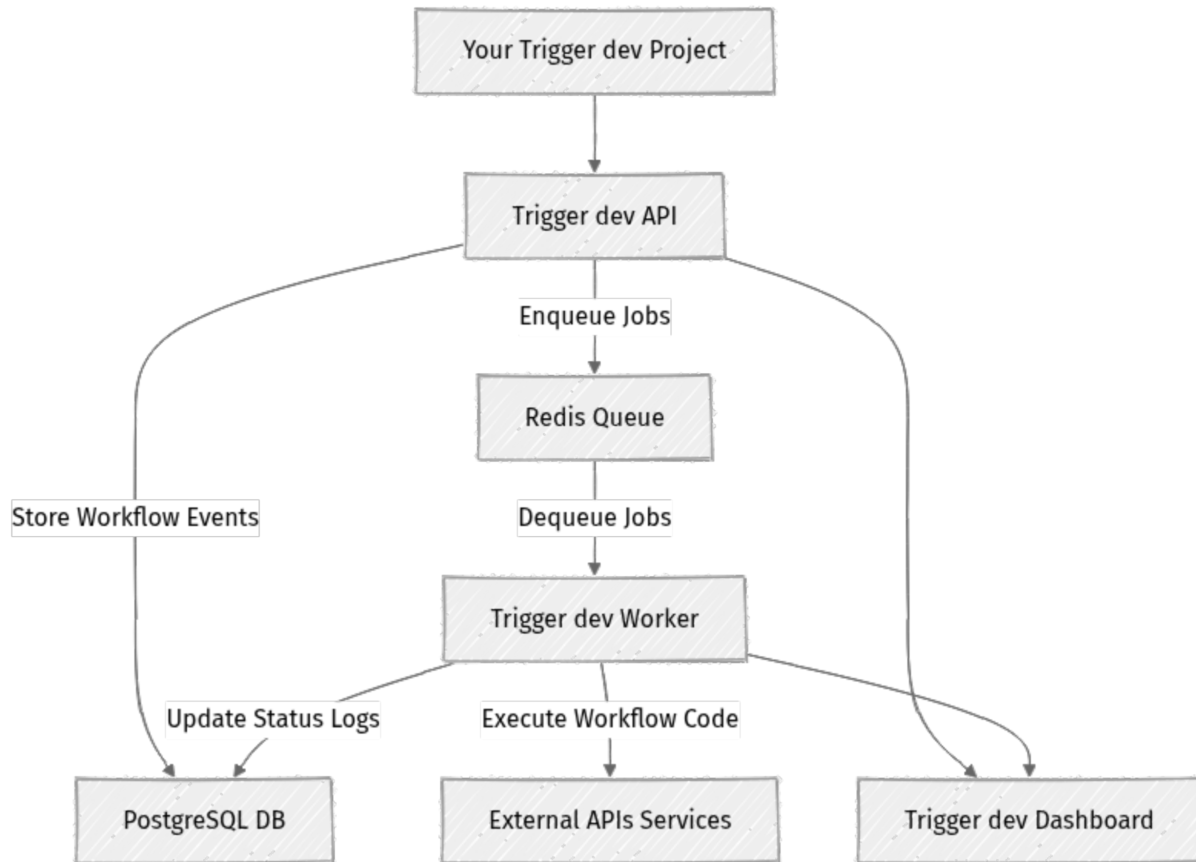

Why Choose Self-Hosting?

The decision to self-host is often driven by specific, advanced needs that the managed service might not fully address:

- **Data Residency & Compliance:** For industries with strict regulations (e.g., healthcare, finance, government), keeping all data within a specific geographic region or on private infrastructure is often a non-negotiable requirement.
- **Deep Customization:** You might need to integrate Trigger.dev with highly specialized internal systems, monitoring tools, or authentication mechanisms that require direct access to the underlying infrastructure and configuration.
- **Cost Optimization at Extreme Scale:** While the managed service offers convenience, for extremely high volumes of workflows and long-term usage, managing your own infrastructure might offer cost savings, provided you have the dedicated DevOps expertise.
- **Enhanced Security Posture:** Some organizations prefer to manage their entire security perimeter, including all infrastructure components, for a unified security strategy and direct control over vulnerability patching.

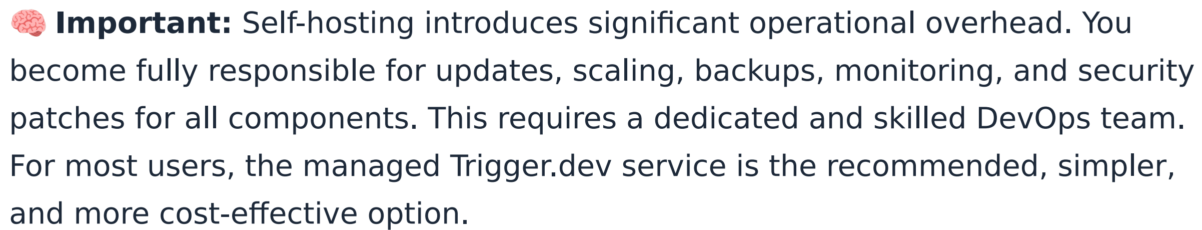

🧠 **Important:** Self-hosting introduces significant operational overhead. You become fully responsible for updates, scaling, backups, monitoring, and security patches for all components. This requires a dedicated and skilled DevOps team. For most users, the managed Trigger.dev service is the recommended, simpler, and more cost-effective option.

Prerequisites for Local Self-Hosting

Before we dive into setting up Trigger.dev locally, ensure you have the following tools and services ready on your development machine.

Local Development Tools

- **Docker Engine (v24.0.0+):** Essential for running containers. Available for Windows, macOS, and Linux.
- **Docker Compose (v2.20.0+):** Simplifies multi-container application management. It's typically bundled with Docker Desktop or installed separately.
- **Node.js (v18.0.0+):** Required for running Trigger.dev client projects and potentially for building specific self-hosted components.
- **Git:** For cloning the Trigger.dev open-source repository.

Infrastructure Components (Provided by Docker Compose Locally)

For a production deployment, you'd provision these as managed cloud services or on your own servers. For local development, our Docker Compose setup will conveniently spin up these instances for you:

- **PostgreSQL Database (v14.0+):** A robust relational database for persistent storage.
- **Redis (v6.0+):** An in-memory data store used for caching and job queuing.

Step-by-Step: Local Self-Hosting with Docker Compose

Let's get Trigger.dev `v4-beta` up and running on your local machine using Docker Compose. This provides a quick and isolated way to test the self-hosted environment before considering production deployments.

Step 1: Clone the Trigger.dev Repository

Your first step is to obtain the Trigger.dev source code from its official GitHub repository. This gives you access to all the necessary Docker Compose files and configuration examples.

```
git clone https://github.com/triggerdotdev/trigger.dev.git
```

This command downloads the entire Trigger.dev project to your current directory.

Step 2: Navigate and Checkout the `v4-beta` Branch

After cloning, navigate into the new `trigger.dev` directory. Then, switch to the `v4-beta` branch. As of 2026-05-20, `v4-beta` represents the cutting edge for upcoming features and the architecture we're interested in for self-hosting.

```
cd trigger.dev
git checkout v4-beta
```

Step 3: Configure Environment Variables

Trigger.dev uses environment variables for configuration, allowing flexible setup without modifying the core code. You'll find an example file in the repository. We'll create a local `.env.local` file to override defaults for our local setup.

Inside the `trigger.dev` directory, copy the example environment file:

```
cp .env.example .env.local
```

Now, open `.env.local` in your favorite text editor. You'll need to set a few critical variables that define how the services connect to each other and how your client projects will connect to your self-hosted instance.

```
# .env.local (simplified for local setup)

# --- Database Configuration ---
# This specifies the connection string for your PostgreSQL database.
# For Docker Compose, 'postgres' is the service name, and '5432' is the
default port.
# 'triggerdev' is the database name.
DATABASE_URL="postgresql://postgres:postgres@localhost:5432/triggerdev?
schema=public"

# --- Redis Configuration ---
# This points to your Redis instance, used for queues and caching.
# 'localhost:6379' is the standard port for Redis.
REDIS_URL="redis://localhost:6379"
```

```
# --- Authentication Secret ---
# IMPORTANT: This secret is used to sign and verify tokens for internal
# communication and client project authentication.
# For production, generate a strong, random, long string!
AUTH_SECRET="super-secret-key-for-local-dev-ONLY"

# --- External API URL (for your client projects to connect to) ---
# This is the URL your Trigger.dev client projects will use to send events
# and register workflows. '2020' is the default port for the API container.
NEXT_PUBLIC_TRIGGER_API_URL="http://localhost:2020"

# --- Dashboard URL (for the dashboard to know its own address) ---
# This is the base URL for the self-hosted dashboard.
# '3000' is the default port for the dashboard container.
NEXT_PUBLIC_APP_ORIGIN="http://localhost:3000"

# --- Other potentially useful variables (check .env.example for full list)
---
# For local development, you might not need to change most others.
# For example, disable telemetry if desired:
TELEMETRY_ENABLED="false"
```

Explanation of Key Variables:

- `DATABASE_URL`: Tells the Trigger.dev API and Worker where to connect to the PostgreSQL database. `postgresql://postgres:postgres@localhost:5432/triggerdev?schema=public` specifies the username, password, host, port, and database name for the local setup.
- `REDIS_URL`: Directs Trigger.dev services to your Redis instance, typically `redis://localhost:6379` for a local Docker Compose setup.
- `AUTH_SECRET`: **This is a critical security credential.** It's used for signing and verifying tokens, ensuring secure communication between Trigger.dev components and authenticating your client projects. **In production, you MUST generate a long, random, cryptographically secure string for this.** The example value is strictly for local development.
- `NEXT_PUBLIC_TRIGGER_API_URL`: This defines the public-facing URL of your self-hosted Trigger.dev API. Your client applications will use this URL to interact with your instance.
- `NEXT_PUBLIC_APP_ORIGIN`: This is the base URL where your self-hosted Trigger.dev Dashboard will be accessible.

Step 4: Start Trigger.dev Services with Docker Compose

With your environment variables configured, you can now spin up all the necessary services (API, Worker, Dashboard, PostgreSQL, Redis) using Docker Compose.

```
docker compose -f docker-compose.yml -f docker-compose.dev.yml up -d
```

Let's break down this command:

- `docker compose`: The command-line tool to interact with Docker Compose.
- `-f docker-compose.yml`: Specifies the main Docker Compose file. This file defines the core Trigger.dev services (API, Worker, Dashboard) along with their dependencies (PostgreSQL and Redis).
- `-f docker-compose.dev.yml`: This option overlays development-specific configurations on top of the main `docker-compose.yml`. For instance, it might expose specific ports to your host machine for local access or enable debug modes.
- `up -d`: This command starts all the services defined in the specified Compose files in "detached" mode, meaning they run in the background without tying up your terminal.

You should see output indicating that Docker is pulling images (if not already cached) and creating/starting containers. This process might take a few minutes the very first time.

To verify that all containers are running as expected, use:

```
docker compose ps
```

You should see entries for `trigger-dev-api`, `trigger-dev-worker`, `trigger-dev-dashboard`, `postgres`, and `redis`, all in a `running` state.

Step 5: Access the Trigger.dev Dashboard

Once all containers are up and running successfully, you can access your self-hosted Trigger.dev dashboard in your web browser:

http://localhost:3000

You'll be prompted to sign up or log in. Since this is a fresh, local instance, create a new account. This account and its associated data will be stored in your local PostgreSQL database.

Step 6: Connect Your Trigger.dev Client Project

Now that your self-hosted Trigger.dev instance is operational, you need to tell your existing Trigger.dev client project (the one where you define your workflows) to connect to this instance instead of the managed cloud service.

In your Trigger.dev client project (e.g., a Next.js app initialized with `npx trigger.dev@v4-beta init`), locate your `.env.local` or `.env` file and update the following environment variables:

```
# In your client project's .env.local

# This must match NEXT_PUBLIC_TRIGGER_API_URL from your self-hosted
instance's .env.local
TRIGGER_API_URL="http://localhost:2020"

# This secret authenticates your client project with your self-hosted API.
# You'll get this from the self-hosted dashboard after creating a project.
TRIGGER_SECRET="YOUR_CLIENT_SECRET_FROM_SELF_HOSTED_DASHBOARD"
```

Explanation:

- `TRIGGER_API_URL`: This environment variable in your client project must match the `NEXT_PUBLIC_TRIGGER_API_URL` you configured in your self-hosted instance's `.env.local` file (e.g., `<http://localhost:2020>`). This directs your client to your local API.
- `TRIGGER_SECRET`: After logging into your self-hosted dashboard (`<http://localhost:3000>`) and creating a new project, you will be provided with a unique `TRIGGER_SECRET` for that project. Copy and paste this secret into your client project's `.env.local` file. This secret authenticates your client project with your self-hosted Trigger.dev API.

Finally, run your client project:

```
npm run dev
```

Your workflows should now register with and execute against your local self-hosted Trigger.dev instance!

Mini-Challenge: Connect a Test Workflow to Local Self-Hosted Instance

It's time to validate your self-hosted setup by deploying and observing a simple workflow.

Challenge:

1. **Verify Setup:** Ensure your local self-hosted Trigger.dev instance is fully running (PostgreSQL, Redis, API, Worker, Dashboard containers should all be active).

2. **Create/Select Client Project:** Use an existing Trigger.dev client project, or create a new one using `npx trigger.dev@v4-beta init`.
3. **Configure Client Project:** Modify your client project's `.env.local` file to correctly connect to your local self-hosted Trigger.dev API and authenticate with the appropriate `TRIGGER_SECRET`.
4. **Implement Simple Workflow:** Create a basic "Hello World" workflow in your client project. This workflow should, for example, log a message to the console and perhaps include a `step.delay(5000)` to simulate a short asynchronous operation.
5. **Trigger Workflow:** Run your client project and trigger the workflow (e.g., via an HTTP endpoint if you set one up, or by manually sending an event).
6. **Observe & Verify:** Navigate to your local self-hosted Trigger.dev dashboard (<http://localhost:3000>) and confirm that the workflow run appears, progresses, and completes successfully, with its logs visible.

Hint:

- Carefully double-check that your client project's `TRIGGER_API_URL` and `TRIGGER_SECRET` exactly match the values you configured for your self-hosted instance.
- Remember to obtain the `TRIGGER_SECRET` from a project created within your self-hosted dashboard, not from the managed Trigger.dev cloud.
- If the workflow doesn't appear, examine the Docker logs for your `trigger-dev-api` and `trigger-dev-worker` containers: `docker compose logs trigger-dev-api -f` and `docker compose logs trigger-dev-worker -f`.

What to Observe/Learn:

- Successful registration of your workflow in the self-hosted dashboard.
- The workflow run appearing, progressing through its steps, and completing, with its associated logs visible in the dashboard.
- This confirms that your client project is correctly communicating with your local self-hosted Trigger.dev API, the API is processing events, and the worker is successfully executing jobs.

Deployment Considerations for Production

Moving from a local Docker Compose setup to a production-grade self-hosted Trigger.dev instance is a significant leap. It involves robust infrastructure, stringent security measures, and meticulous operational best practices. This section outlines key considerations for a reliable and scalable production deployment.

⚠️ **What can go wrong:** The self-hosting guide for Trigger.dev `v4-beta` is still an active area of development, as indicated by open GitHub issues like #48 and #2186 (checked 2026-05-20). This means the process can be complex, and official, comprehensive documentation for production deployments might still be evolving. Proceed with caution, refer to the official GitHub repository for the latest guidance, and be prepared for potential adjustments.

1. Scalability and High Availability

- **Load Balancing for API:** Deploy multiple instances of the Trigger.dev API behind a robust load balancer (e.g., AWS ALB, Nginx, HAProxy). This distributes incoming traffic, prevents single points of failure, and allows for zero-downtime updates.
- **Auto-Scaling Workers:** Implement auto-scaling for your Trigger.dev Worker instances. Monitor metrics like Redis queue depth or worker CPU utilization to automatically adjust the number of workers, ensuring your workflows can handle varying loads efficiently.
- **Highly Available Database:** For PostgreSQL, move beyond a single instance. Consider a highly available setup using:
  - **Managed Services:** AWS RDS, Azure Database for PostgreSQL, Google Cloud SQL, which handle replication, failover, and backups automatically.
  - **Clustering Solutions:** Patroni, PgBouncer, or manual replication for self-managed clusters, providing read replicas and automatic failover.
- **Redis Persistence & Clustering:** Configure Redis for persistence (using AOF or RDB snapshots) to prevent data loss. For high availability and scalability, consider a Redis Cluster or Sentinel setup.

2. Monitoring and Logging

- **Centralized Logging:** Aggregate all Trigger.dev component logs (API, Worker, Dashboard, PostgreSQL, Redis) into a centralized logging solution. Popular choices include Elastic Stack (ELK), Grafana Loki, or cloud-native services like AWS CloudWatch Logs, Azure Monitor, or Google Cloud Logging. This is crucial for debugging distributed workflows.
- **Metrics & Alerts:** Collect detailed metrics (CPU, memory, network I/O, disk usage, queue depth, job counts, workflow durations) from all components. Use monitoring tools like Prometheus for collection and Grafana for visualization. Set up automated alerts for critical thresholds or service failures.
- **Distributed Tracing:** For complex, long-running workflows that interact with multiple external services, integrate a distributed tracing system (e.g., OpenTelemetry with Jaeger or Zipkin). This provides end-to-end visibility into the execution path and latency of requests across all services.

3. Security Best Practices

- **Network Isolation:** Deploy all Trigger.dev components within a private network (VPC). Implement strict firewall rules (Security Groups, Network ACLs) to limit ingress and egress traffic, exposing only necessary ports (e.g., HTTPS for API/Dashboard) to the public internet.
- **Secrets Management:** Never hardcode sensitive credentials. Use a dedicated secrets management service (e.g., AWS Secrets Manager, Azure Key Vault, HashiCorp Vault, Kubernetes Secrets) for `AUTH_SECRET`, database credentials, API keys, and other sensitive configuration.
- **Access Control & Authentication:** Implement strong authentication and authorization for the Trigger.dev Dashboard and API. If possible, integrate with your organization's existing identity provider (e.g., SSO).
- **Regular Updates:** Keep all components, including the underlying operating system, Docker, PostgreSQL, Redis, and Trigger.dev application code, regularly updated to patch security vulnerabilities and benefit from bug fixes.

4. Deployment Strategy

- **Container Orchestration:** For production, Docker Compose is a development tool. Use a robust container orchestration platform like Kubernetes (managed services like EKS, AKS, GKE are ideal) or Docker Swarm for managing, scaling, and self-healing your containers.

- **CI/CD Pipelines:** Automate the entire deployment process using Continuous Integration/Continuous Deployment (CI/CD) pipelines (e.g., GitHub Actions, GitLab CI, Jenkins). This ensures consistent, reliable, and repeatable deployments.
- **Database Migrations:** Trigger.dev includes database schema migration scripts. Implement a strategy to apply these migrations safely as part of your deployment process, ideally with zero downtime.

5. Backup and Disaster Recovery

- **Database Backups:** Implement robust and regular backup procedures for your PostgreSQL database. Ensure backups are stored off-site and test your recovery process periodically to verify data integrity and recovery time objectives (RTO).
- **Infrastructure as Code (IaC):** Define your entire Trigger.dev infrastructure (VMs, networks, databases, load balancers) using Infrastructure as Code tools like Terraform or Pulumi. This enables consistent environment provisioning and quick recovery in case of a disaster.

Common Pitfalls & Troubleshooting Self-Hosting

Self-hosting complex distributed systems like Trigger.dev can introduce unique challenges. Here are some common pitfalls and effective troubleshooting strategies, including some specific to Trigger.dev's architecture.

1. Incorrect Environment Variables

- **Symptom:** Services fail to start, the API returns authentication errors (e.g., 401 Unauthorized), or the dashboard cannot connect to the API.

- **Troubleshooting:**
    - **Review `.env.local`:** Carefully compare your `.env.local` (or production environment variables) against the `.env.example` in the Trigger.dev repository. Pay close attention to typos or missing variables.
    - **`AUTH_SECRET`:** Ensure `AUTH_SECRET` is set and is a strong, unique string. Mismatches or weak secrets are a common cause of authentication failures.
    - **Database/Redis URLs:** Verify `DATABASE_URL` and `REDIS_URL` are correctly formatted and point to accessible instances.
    - **Client Project Variables:** Double-check that your client project's `TRIGGER_API_URL` and `TRIGGER_SECRET` accurately match your self-hosted instance's configuration.

2. Port Conflicts

- **Symptom:** Docker containers fail to start with "port already in use" errors.
- **Troubleshooting:**
    - **Check Port Usage:** Identify if other applications on your host machine are using ports `2020` (Trigger.dev API), `3000` (Trigger.dev Dashboard), `5432` (PostgreSQL), or `6379` (Redis).
    - **Modify `docker-compose.dev.yml`:** If conflicts exist, you can modify the exposed ports in `docker-compose.dev.yml` (under the `ports` section for each service). Remember to update `NEXT_PUBLIC_TRIGGER_API_URL` and `NEXT_PUBLIC_APP_ORIGIN` in your `.env.local` file accordingly.

3. Database Connectivity Issues

- **Symptom:** Trigger.dev API or Worker containers crash on startup with database connection errors (e.g., `FATAL: database "triggerdev" does not exist`, `connection refused`).

- **Troubleshooting:**
    - **PostgreSQL Status:** Verify that your PostgreSQL container (or external instance) is running and accessible (`docker compose ps`).
    - **Firewall Rules:** If PostgreSQL is on a different host or network, check firewall rules to ensure the Trigger.dev containers can reach it.
    - **Credentials:** Confirm that the username, password, and database name in `DATABASE_URL` are correct.
    - **Database Creation:** Ensure the `triggerdev` database exists or that the Trigger.dev services have permissions to create it on first run (which Docker Compose typically handles).

4. Redis Connection Problems

- **Symptom:** Workflows get stuck, jobs aren't processed, the dashboard feels unresponsive, or events are not triggering workflows.
- **Troubleshooting:**
    - **Redis Status:** Verify your Redis container (or external instance) is running (`docker compose ps`).
    - **`REDIS_URL`:** Double-check the `REDIS_URL` in your environment variables for correctness.
    - **Resource Usage:** Ensure Redis isn't running out of memory or connections. Monitor Redis metrics.

5. Resource Exhaustion

- **Symptom:** Services are slow, unresponsive, or crash unpredictably, especially under load.
- **Troubleshooting:**
    - **Monitor Resources:** Use `docker stats` for local Docker containers or cloud monitoring tools for production VMs/containers to track CPU, memory, and disk I/O usage.
    - **Allocate More:** Increase the allocated resources (CPU, RAM) to your Docker daemon, host machine, or cloud compute instances.

6. Outdated Codebase / Schema Mismatches

- **Symptom:** Unexpected behavior, missing features, or compatibility issues between Trigger.dev services and your client project, or database errors after an update.

- **Troubleshooting:**
    - **`git pull`:** Regularly `git pull` from the `v4-beta` branch of the Trigger.dev repository to get the latest code.
    - **Rebuild Images:** After pulling new code, rebuild your Docker images: `docker compose build`.
    - **Database Migrations:** If the database schema has changed, Trigger.dev typically includes migration scripts. Ensure these are run as part of your update process. You might need to manually trigger them if they don't run automatically on container startup.

7. State Management in Long-Running Workflows

- **Symptom:** Workflow state seems to reset, or data is missing between `waitFor` or `delay` steps, leading to incorrect execution paths or data loss.
- **Troubleshooting:**
    - **Inspect Database:** Connect directly to your PostgreSQL database. Examine the `Run` and `Job` tables for the specific workflow run. Pay close attention to the `payload` and `state` columns, which store the workflow's context and durable state.
    - **Worker Logs:** Check your `trigger-dev-worker` logs for any errors related to database writes, deserialization failures, or issues when persisting/retrieving state from PostgreSQL.
    - **Workflow Code Logic:** Review your workflow code carefully. Ensure that state is being correctly passed between steps, especially when using `step.waitFor`, `step.delay`, or other durable execution primitives that rely on state serialization and deserialization.

8. Debugging Distributed Trigger.dev Workflows

- **Symptom:** A workflow fails, but the relevant logs are scattered across multiple API or Worker instances, making it extremely difficult to trace the execution path and pinpoint the root cause.

- **Troubleshooting:**
    - **Centralized Logging (Critical):** This is paramount for distributed systems. Implement a robust centralized logging solution (e.g., ELK Stack, Grafana Loki, or cloud-native logging services) to aggregate logs from all Trigger.dev API and Worker instances into a single searchable location.
    - **Correlation IDs:** Leverage the workflow run ID provided by Trigger.dev. This ID acts as a correlation token. Use it to filter and search logs across all services to trace a single execution from start to finish.
    - **Distributed Tracing (Advanced):** Integrate an OpenTelemetry-based system for end-to-end tracing. This visually maps out the calls between services and provides detailed timing information, making it easy to identify bottlenecks or failure points.
    - **Trigger.dev Dashboard:** The self-hosted dashboard (`<http://localhost:3000>`) is your primary source of truth. It aggregates logs for specific runs, shows detailed status, and provides error details, significantly simplifying initial debugging.

9. Event Ingestion and Processing Issues

- **Symptom:** Events sent to Trigger.dev are not triggering workflows, or there's a significant, unexpected delay in processing them.
- **Troubleshooting:**
    - **API Connectivity:** Verify that your client application or event source can successfully reach your self-hosted `NEXT_PUBLIC_TRIGGER_API_URL`. Check network firewalls, DNS resolution, and any proxies in the path.
    - **Event Matchers:** Double-check your workflow's `trigger` definition. Ensure that the event payload structure and properties precisely match what your workflow is expecting. Mismatches are a very common cause of untriggered workflows.
    - **Redis Queue Health:** Monitor the depth of your Redis queues. A rapidly growing queue suggests that your workers are not processing events fast enough, or there's an issue with worker connectivity to Redis.
    - **API Logs:** Inspect the logs of your `trigger-dev-api` containers for any errors during event ingestion, validation, or routing processes.

10. Worker Idempotency Failures

- **Symptom:** External actions (e.g., sending an email, writing to a database, calling a third-party API) are duplicated, even though the workflow step should logically only run once.
- **Troubleshooting:**
    - **Design for Idempotency:** The most robust solution is to design the external services your workflows interact with to be inherently idempotent. This means that calling an operation multiple times with the same input has the same effect as calling it once.
    - **Trigger.dev Retries:** Understand that Trigger.dev's durable execution includes automatic retries for failed steps. If an external call succeeds but the acknowledgment back to Trigger.dev fails (e.g., network glitch), the step might be retried. Your external service must handle this.
    - **Examine Run History:** Use the Trigger.dev dashboard to inspect the detailed run history. Look for steps that were marked as retried or executed multiple times, which can indicate an idempotency issue.

11. General Debugging Strategy

- **Docker Logs First:** Always start by checking the Docker logs for the relevant service (`docker compose logs <service_name>`). Use `-f` to follow logs in real-time.
- **Trigger.dev Dashboard:** This is your central hub for workflow-specific debugging. It provides aggregated logs, status, and error details for each run.
- **Database Inspection:** For persistent state issues, connect directly to your PostgreSQL database and inspect tables like `Job`, `Run`, and `Event` to understand the system's internal state.

Summary

In this chapter, we embarked on an advanced journey into self-hosting Trigger.dev. This path offers unparalleled control and customization, but also demands significant operational responsibility.

Here are the key takeaways from our exploration:

- **Why Self-Host:** The decision to self-host is typically driven by specific needs such as data residency, deep infrastructure customization, potential cost optimization at extreme scale, or enhanced security posture.
- **Core Architecture:** Trigger.dev is a distributed system comprising several interconnected components: the API, Worker, Dashboard, PostgreSQL (for durable state), and Redis (for queuing and caching).
- **Local Setup:** We successfully set up a local self-hosted Trigger.dev instance using Docker Compose, configuring essential environment variables, and connecting a client project to it.
- **Production Readiness:** Deploying to a production environment requires careful consideration of scalability (load balancing, auto-scaling), high availability (database clustering, Redis persistence), robust monitoring and logging, stringent security practices, and a mature CI/CD deployment strategy.
- **Ongoing Development:** It's crucial to acknowledge that self-hosting for Trigger.dev `v4-beta` is an evolving area. Staying updated with the official GitHub repository for the latest guidance is essential.
- **Troubleshooting:** Be prepared to debug common issues such as incorrect environment variables, port conflicts, database/Redis connectivity, resource exhaustion, and Trigger.dev-specific challenges like durable state management, distributed workflow debugging, event ingestion issues, and worker idempotency failures.

Self-hosting Trigger.dev is a powerful option for organizations and advanced developers who require deep control over their workflow infrastructure. While complex, mastering this setup empowers you to tailor Trigger.dev to the most demanding environments and integrate it seamlessly into existing enterprise systems.

References

- Trigger.dev GitHub Repository: https://github.com/triggerdotdev/trigger.dev
- Trigger.dev Documentation: https://trigger.dev
- Docker Compose Overview: https://docs.docker.com/compose/
- PostgreSQL Official Website: https://www.postgresql.org/
- Redis Official Website: https://redis.io/

This page is AI-assisted and reviewed. It references official documentation and recognized resources where relevant.

CHAPTER 11

Real-World Project: Building an AI-Powered Customer Support Agent

Building intelligent automation often means dealing with complex, multi-step processes that might involve external services, human intervention, and unpredictable delays. This is especially true for AI agents that interact with users and critical systems.

In this chapter, we'll put all our Trigger.dev knowledge to the test by creating a practical, real-world AI-powered customer support agent. You'll learn how to orchestrate an AI agent workflow that can classify user queries, retrieve information from a knowledge base, and even escalate to a human agent when needed, all while maintaining state across long-running, durable executions.

This project will solidify your understanding of Trigger.dev's core features like AI agents, durable execution, human-in-the-loop workflows, and robust error handling. To get the most out of this chapter, you should be comfortable with basic Trigger.dev workflow setup, understand `io.run` and `io.wait`, and have a foundational grasp of asynchronous JavaScript/TypeScript.

Core Concepts: Architecting an Intelligent Support System

An effective AI customer support agent isn't just a chatbot; it's a sophisticated system capable of understanding intent, accessing information, and making decisions, often involving human oversight. Trigger.dev provides the perfect foundation for building such a system. Let's break down the key architectural considerations.

The Role of AI in Customer Support

At its heart, an AI support agent aims to solve customer problems efficiently. This involves several critical capabilities:

1. **Understand User Intent:** What is the customer trying to achieve or ask? AI models excel at discerning the underlying goal from natural language.
2. **Provide Relevant Information:** Answer questions directly from a knowledge base or by synthesizing information from various sources.

3. **Automate Common Tasks:** Handle routine requests like checking order status or updating personal details without human intervention, freeing up human agents.
4. **Escalate When Necessary:** Identify situations too complex, ambiguous, or sensitive for AI and seamlessly hand them off to a human expert.

Human-in-the-Loop (HITL) for Trust and Accuracy

While AI is powerful, it's not infallible. For critical applications like customer support, human oversight is essential. Human-in-the-Loop (HITL) workflows ensure that complex, ambiguous, or high-stakes decisions are reviewed and approved by a human.

Trigger.dev excels at HITL by allowing workflows to pause and `wait` for external events. For instance, a workflow can wait for a human to approve an AI-generated response or provide missing input. This creates robust systems where AI handles the mundane, and humans focus on value-added tasks that require empathy, nuanced judgment, or creative problem-solving.

Durable Workflows for Conversational AI

Customer interactions are rarely a single request-response. They often involve multi-turn conversations, where context from previous messages is crucial. Trigger.dev's durable execution is a game-changer here:

- **State Persistence:** The workflow's state (e.g., conversation history, user ID, current query context) is automatically saved and restored. This means if the underlying server crashes or restarts, the conversation can pick up exactly where it left off, ensuring a seamless experience for the customer.
- **Long-Running Operations:** Workflows can pause for minutes, hours, or even days (e.g., while waiting for human input or an external API response) without consuming active server resources. They resume precisely at the point of interruption.
- **Retries and Observability:** If an external API call fails (e.g., to an AI service or a CRM), Trigger.dev automatically retries according to configured policies. You also get clear observability into each step of the conversation directly from the Trigger.dev dashboard.

Agent Workflow Diagram

Let's visualize the flow of our AI customer support agent. This diagram illustrates the decision points and interactions between AI and human agents.

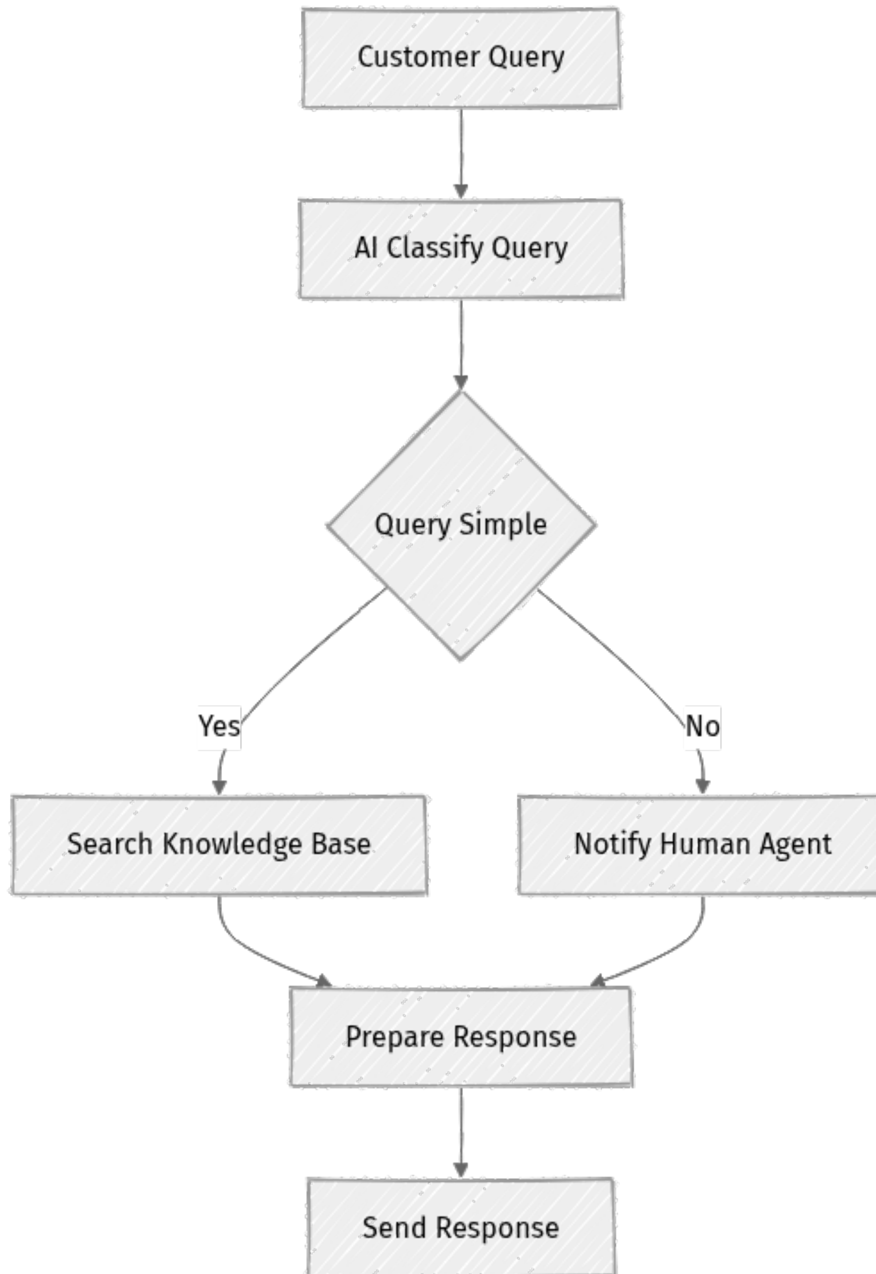

This diagram shows how a customer query initiates a Trigger.dev workflow. The AI first classifies the query. Simple queries are handled by searching a knowledge base, while complex ones trigger a human escalation. Both paths eventually lead to a response being sent back to the customer, potentially after human review.

Step-by-Step Implementation: Building the Agent

We'll build this agent incrementally, focusing on clarity and understanding each piece as we add it.

1. Project Setup

First, let's create a new Trigger.dev project. We'll be using the `v4-beta`, which is stable and feature-rich, with General Availability (GA) expected around May/June 2026. This version includes many enhancements for AI workflows.

Open your terminal and run the following command:

```
npx trigger.dev@v4-beta init
```

Follow the prompts:

- `What is your project name?` Enter `ai-support-agent`.
- `Which framework are you using?` Choose `Next.js` (or your preferred framework; the core Trigger.dev logic remains similar).
- `Do you want to use TypeScript?` `Yes`.
- `Do you want to install dependencies?` `Yes`.

Once the project is created and dependencies are installed, navigate into your new project directory:

```
cd ai-support-agent
```

We'll also need an AI SDK, like OpenAI's, and a simple way to simulate external services. We'll use `zod` for robust schema validation, which is crucial when dealing with AI outputs.

```
npm install openai@4.x zod
```

`zod` is a schema declaration and validation library that is very useful for defining the structure of our AI responses or human input. `openai@4.x` is the latest major version as of 2026-05-20.

2. Configure Trigger.dev Client and Environment Variables

You'll need your Trigger.dev API key to connect your application to the Trigger.dev cloud. In your `.env` file, add these variables:

```
# .env
TRIGGER_API_KEY="your_trigger_api_key_here"
OPENAI_API_KEY="your_openai_api_key_here"
```

Replace `your_trigger_api_key_here` with your actual key from the Trigger.dev dashboard. Similarly, replace `your_openai_api_key_here` with your OpenAI API key.

Next, open `src/trigger.ts` (or `app/trigger.ts` depending on your Next.js setup) and ensure your client is initialized correctly, referencing the API key from the environment.

```
// src/trigger.ts (or app/trigger.ts)
import { TriggerClient } from "@trigger.dev/sdk";

export const client = new TriggerClient({
  id: "ai-support-agent",
  apiKey: process.env.TRIGGER_API_KEY,
  apiUrl: process.env.TRIGGER_API_URL, // Optional: defaults to
cloud.trigger.dev for the managed service
});
```

3. Define the AI Agent Workflow

Now, let's create our main workflow. We'll define it in a new file, `src/jobs/customerSupportAgent.ts`. This file will contain the entire logic for our AI agent.

Step 1: Initial Workflow Structure and AI Classification

We'll start by defining the job and using an AI model to classify the incoming customer query. For simplicity, we'll use the actual `openai` package for AI calls.

Create `src/jobs/customerSupportAgent.ts` and add the following code:

```
// src/jobs/customerSupportAgent.ts
import { client } from "@/trigger"; // Adjust path as needed for your project
structure
import { OpenAI } from "openai"; // For actual OpenAI API calls
import { z } from "zod"; // For schema validation

// Initialize OpenAI client
const openai = new OpenAI({
  apiKey: process.env.OPENAI_API_KEY,
});

// Define the schema for the AI's classification response
const ClassificationSchema = z.object({
  type: z.enum(["simple_faq", "complex_issue", "billing_query", "feature_reque
st", "unknown"]),
  keywords: z.array(z.string()).optional(),
  confidence: z.number().min(0).max(1),
  summary: z.string().optional(),
});

client.defineJob({
  id: "customer-support-agent",
  name: "AI Customer Support Agent",
  version: "1.0.0",
```

```
  enabled: true,
  trigger: client.defineEventTrigger({
    name: "customer.query",
    schema: z.object({
      userId: z.string(),
      query: z.string(),
      conversationId: z.string().optional(),
    }),
  }),
  run: async (payload, io, ctx) => {
    io.logger.info("Received customer query", payload);

    // 📌 Key Idea: Use io.runTask for external API calls like AI inference.
    // It provides automatic retries, observability in the dashboard, and
error handling.
    const classification = await io.runTask(
      "classify-query",
      {
        name: "Classify Customer Query with AI",
        icon: "openai", // Use a relevant icon for dashboard visibility
        params: { query: payload.query },
      },
      async (task, params) => {
        // Here, we're calling the OpenAI API to classify the query.
        const response = await openai.chat.completions.create({
          model: "gpt-4o", // Using the latest capable model as of 2026-05-20
          messages: [
            { role: "system", content: `You are a helpful assistant that
classifies customer queries into one of these categories: simple_faq,
complex_issue, billing_query, feature_request, or unknown. Provide relevant
keywords, a confidence score (0-1), and a brief summary. Respond strictly in
JSON format.` },
            { role: "user", content: `Classify the following query: "$
{params.query}"` },
          ],
          response_format: { type: "json_object" },
          temperature: 0.1, // Lower temperature for more deterministic output
        });

        const rawResult = response.choices[0].message?.content;
        if (!rawResult) {
          throw new Error("AI classification failed to return content.");
        }
        const parsedResult = JSON.parse(rawResult);

        // 🧠 Important: Validate AI output using Zod for robustness.
        // This ensures the AI's response matches our expected structure.
        const validatedClassification = ClassificationSchema.parse(parsedResul
t);
        return validatedClassification;
      }
    );

    io.logger.info("Query classified", { classification });

    // For now, let's just log the classification. We'll add more logic next.
    return { status: "classified", classification };
  },
});

// Helper function to simulate knowledge base lookup
// ⚡ Real-world insight: This would typically hit a real database, search
```

```
index,
// or a dedicated knowledge base API (e.g., Zendesk, Salesforce).
async function fetchKnowledgeBase(query: string): Promise<string | null> {
  const articles = {
    "shipping status":
"Your order usually ships within 2-3 business days. You can track it here:
[tracking link].",
    "refund policy": "Our refund policy allows returns within 30 days of
purchase, provided the item is unused and in original packaging.",
    "account login": "If you're having trouble logging in, please try
resetting your password. If that doesn't work, contact support.",
    "password reset": "To reset your password, visit our password recovery
page and follow the instructions.",
  };
  const lowerQuery = query.toLowerCase();
  for (const keyword in articles) {
    if (lowerQuery.includes(keyword)) {
      return articles[keyword];
    }
  }
  return null;
}
```

Explanation:

- We import `TriggerClient`, the `OpenAI` SDK, and `zod`.
- An `OpenAI` client is initialized using the `OPENAI_API_KEY` from your environment.
- `ClassificationSchema` uses Zod to define the expected structure of the AI's output. This is crucial for ensuring our workflow only proceeds with valid and predictable data from the AI.
- `client.defineJob` sets up our `customer-support-agent` job. It's configured to trigger on a custom event named `customer.query`, which includes `userId`, `query`, and an optional `conversationId`.
- Inside the `run` function, `io.runTask` is used to execute the AI classification. This ensures that the AI call is retried if it fails, and its execution (inputs, outputs, logs) is clearly visible in the Trigger.dev dashboard.
- The AI prompt instructs `gpt-4o` to classify the query into predefined categories and respond in JSON. We then parse this JSON and validate it with Zod.
- A `fetchKnowledgeBase` mock function is included, ready for later use to simulate retrieving answers from a knowledge base.

To test this initial setup, save the file, then run your Next.js development server:

```
npm run dev
```

In a separate terminal, trigger the `customer.query` event:

```
npx trigger.dev@v4-beta send customer.query '{"userId": "user-123", "query":
"What is your refund policy?"}'
```

You should see the job run in your Trigger.dev dashboard. Navigate to the dashboard, find the execution, and observe the "Classify Customer Query with AI" step, complete with its input, output, and logs.

Step 2: Conditional Logic and Knowledge Base Lookup

Now, let's add the logic to handle different classification types. If the AI classifies a query as `simple_faq` with high confidence, we'll attempt to answer it from our mock knowledge base.

Modify the `run` function in `src/jobs/customerSupportAgent.ts` by adding the following `if` block after the `io.logger.info("Query classified", { classification });` line:

```
// src/jobs/customerSupportAgent.ts (continued)
// ... (previous code for imports, client, OpenAI, ClassificationSchema,
fetchKnowledgeBase)

client.defineJob({
  id: "customer-support-agent",
  name: "AI Customer Support Agent",
  version: "1.0.0",
  enabled: true,
  trigger: client.defineEventTrigger({
    name: "customer.query",
    schema: z.object({
      userId: z.string(),
      query: z.string(),
      conversationId: z.string().optional(),
    }),
  }),
  run: async (payload, io, ctx) => {
    io.logger.info("Received customer query", payload);

    const classification = await io.runTask(
      "classify-query",
      {
        name: "Classify Customer Query with AI",
        icon: "openai",
        params: { query: payload.query },
      },
      async (task, params) => {
        const response = await openai.chat.completions.create({
          model: "gpt-4o",
          messages: [
            { role: "system", content: `You are a helpful assistant that
```

```
classifies customer queries into one of these categories: simple_faq,
complex_issue, billing_query, feature_request, or unknown. Provide relevant
keywords, a confidence score (0-1), and a brief summary. Respond strictly in
JSON format.` },
            { role: "user", content: `Classify the following query: "$
{params.query}"` },
          ],
          response_format: { type: "json_object" },
          temperature: 0.1,
        });

        const rawResult = response.choices[0].message?.content;
        if (!rawResult) {
          throw new Error("AI classification failed to return content.");
        }
        const parsedResult = JSON.parse(rawResult);
        const validatedClassification = ClassificationSchema.parse(parsedResul
t);
        return validatedClassification;
      }
    );

    io.logger.info("Query classified", { classification });

    let responseToCustomer: string | null = null; // Initialize a variable to
hold the final response

    // 🧠 Important: Conditional logic based on AI classification results.
    if (classification.type === "simple_faq" && classification.confidence > 0.
7) {
      io.logger.info("Attempting to answer simple FAQ from knowledge base.");
      const kbAnswer = await io.runTask(
        "knowledge-base-lookup",
        {
          name: "Lookup Answer in Knowledge Base",
          icon: "database", // Use a database icon for this task
          params: { query: payload.query },
        },
        async (task, params) => {
          return await fetchKnowledgeBase(params.query);
        }
      );

      if (kbAnswer) {
        responseToCustomer = kbAnswer;
        io.logger.info("Found answer in KB.", { answer: kbAnswer });
      } else {
        io.logger.warn("Could not find answer in KB, will consider human
escalation.");
        // If KB fails, responseToCustomer remains null, leading to escalation
later
      }
    }

    // Next, we'll add human escalation if no automated response was found.
    // ... (rest of the run function will go here)
  },
});
```

Explanation:

- We introduce a `responseToCustomer` variable, initialized to `null`. This variable will hold the answer if one is found, whether by AI or human.
- An `if` statement checks if the AI classified the query as `simple_faq` with a confidence score above `0.7`. This threshold helps filter out less certain AI answers.
- If the conditions are met, `io.runTask` calls our `fetchKnowledgeBase` helper function. This task is also retriable and observable.
- If an answer is found in the knowledge base, it's assigned to `responseToCustomer`. If not, `responseToCustomer` remains `null`, signaling that further action (like human escalation) might be needed.

Test this by sending a simple query that should be in our mock knowledge base:

```
npx trigger.dev@v4-beta send customer.query '{"userId": "user-123", "query": "What is the refund policy?"}'
```

Then, send a query that won't be found:

```
npx trigger.dev@v4-beta send customer.query '{"userId": "user-123", "query": "How do I connect my smart fridge?"}'
```

In the Trigger.dev dashboard, you should observe the "Lookup Answer in Knowledge Base" task. For the second query, it will likely complete without finding an answer, and the workflow will terminate (for now) after logging the warning.

Step 3: Human Escalation (Human-in-the-Loop)

What if the AI can't confidently answer a query, or if the query is inherently complex and requires human judgment? This is where human intervention, or Human-in-the-Loop (HITL), comes into play.

Continue modifying the `run` function in `src/jobs/customerSupportAgent.ts` by adding the following `if (!responseToCustomer)` block after the previous `if` statement:

```
// src/jobs/customerSupportAgent.ts (continued)
// ... (previous code for imports, client, OpenAI, ClassificationSchema,
fetchKnowledgeBase)

client.defineJob({
  id: "customer-support-agent",
  name: "AI Customer Support Agent",
  version: "1.0.0",
```

```
  enabled: true,
  trigger: client.defineEventTrigger({
    name: "customer.query",
    schema: z.object({
      userId: z.string(),
      query: z.string(),
      conversationId: z.string().optional(),
    }),
  }),
  run: async (payload, io, ctx) => {
    io.logger.info("Received customer query", payload);

    const classification = await io.runTask(
      "classify-query",
      {
        name: "Classify Customer Query with AI",
        icon: "openai",
        params: { query: payload.query },
      },
      async (task, params) => {
        const response = await openai.chat.completions.create({
          model: "gpt-4o",
          messages: [
            { role: "system", content: `You are a helpful assistant that
classifies customer queries into one of these categories: simple_faq,
complex_issue, billing_query, feature_request, or unknown. Provide relevant
keywords, a confidence score (0-1), and a brief summary. Respond strictly in
JSON format.` },
            { role: "user", content: `Classify the following query: "$
{params.query}"` },
          ],
          response_format: { type: "json_object" },
          temperature: 0.1,
        });

        const rawResult = response.choices[0].message?.content;
        if (!rawResult) {
          throw new Error("AI classification failed to return content.");
        }
        const parsedResult = JSON.parse(rawResult);
        const validatedClassification = ClassificationSchema.parse(parsedResul
t);
        return validatedClassification;
      }
    );

    io.logger.info("Query classified", { classification });

    let responseToCustomer: string | null = null;
    let humanProvidedResponse = false; // Track if a human provided the final
response

    if (classification.type === "simple_faq" && classification.confidence > 0.
7) {
      io.logger.info("Attempting to answer simple FAQ from knowledge base.");
      const kbAnswer = await io.runTask(
        "knowledge-base-lookup",
        {
          name: "Lookup Answer in Knowledge Base",
          icon: "database",
          params: { query: payload.query },
        },
```

```
        async (task, params) => {
          return await fetchKnowledgeBase(params.query);
        }
      );

      if (kbAnswer) {
        responseToCustomer = kbAnswer;
        io.logger.info("Found answer in KB.", { answer: kbAnswer });
      } else {
        io.logger.warn("Could not find answer in KB, will consider human
escalation.");
      }
    }

    // ⚠️ What can go wrong: If no automated response is found,
    // or if the query is complex (e.g., classification.type !==
"simple_faq"),
    // we need human intervention.
    if (!responseToCustomer) {
      io.logger.info("Escalating to human agent for review.");

      // Simulate sending a notification to a human agent dashboard/email.
      // This task would typically integrate with external communication
systems.
      await io.runTask(
        "notify-human-agent",
        {
          name: "Notify Human Agent of Escalation",
          icon: "bell", // Use a bell icon for notifications
          params: {
            userId: payload.userId,
            query: payload.query,
            classification: classification.summary || classification.type,
            runId: ctx.run.id, // Pass the run ID so the human can reference
it
          },
        },
        async (task, params) => {
          // In a real system, this would send an email, Slack message,
          // or create a ticket in a CRM like Zendesk or Salesforce.
          io.logger.warn(`Human Agent NOTIFIED: User ${params.userId} needs
help with: ${params.query}. Run ID: ${params.runId}`);
          // Return a URL for the human to interact with this specific case.
          return `https://your-agent-dashboard.com/review/${ctx.run.id}`;
        }
      );

      // ⚡ Quick Note: io.wait is crucial for Human-in-the-Loop workflows.
      // It pauses the workflow indefinitely (up to the timeout) until an
      // external event is received, without consuming active compute
resources.
      const humanInput = await io.wait(
        "wait-for-human-input",
        {
          name: "Wait for Human Agent Input",
          timeoutInSeconds: 60 * 60 * 24 * 7, // Wait for 7 days for human
input
          // Define the schema for the event we are waiting for.
          // This event would typically be sent from your human agent
dashboard
          // when they submit their review or response.
          // Example: client.sendEvent("human.input", { runId: ctx.run.id,
```

```
response: "...", approved: true });
            schema: z.object({
              response: z.string().min(1, "Response cannot be empty."),
              approved: z.boolean().optional().default(true),
            }),
          }
        );

        io.logger.info("Human agent provided input.", { humanInput });
        if (humanInput.approved) {
          responseToCustomer = humanInput.response;
          humanProvidedResponse = true;
        } else {
          // Human agent disapproved, maybe they need more info or
          // it requires a different action. For this example, we'll
          // still use their response but log the disapproval.
          responseToCustomer = "The human agent reviewed your query. " + humanIn
put.response;
          io.logger.warn("Human agent disapproved the automated response, but
provided an alternative.");
        }
      }

    // ⚡ Real-world insight: If after all steps, we still don't have a
response,
    // it's a critical failure or needs direct human contact.
    if (!responseToCustomer) {
      io.logger.error("Failed to generate a response, even after human
escalation.");
      // In a real system, you might create a high-priority ticket or directly
call the user.
      throw new Error("Unable to resolve customer query automatically or with
human input.");
    }

    // Finally, send the response to the customer.
    // ... (rest of the run function will go here)
  },
});
```

Explanation:

- We introduce a `humanProvidedResponse` boolean to track whether the final response originated from a human agent.
- The `if (!responseToCustomer)` block handles the escalation. This condition triggers if the AI and knowledge base couldn't provide a satisfactory answer.
- `io.runTask("notify-human-agent", ...)` simulates notifying a human agent. In a production system, this would trigger an email, a Slack message, or create a task in a dedicated agent dashboard. The task returns a URL for the human to access the specific case.

- `io.wait("wait-for-human-input", ...)` is the core of the HITL. It pauses the workflow at this exact point.
    - The `timeoutInSeconds` ensures the workflow doesn't wait forever, providing a safety net.
    - The `schema` for `io.wait` defines the expected structure of the external event that will resume this workflow. A human agent would typically interact with a UI that, upon submission, sends an event back to Trigger.dev (e.g., `client.sendEvent("human.input", { runId: ctx.run.id, response: "..." })`) matching this schema.
- Once the human input event is received, the workflow resumes, and `responseToCustomer` is updated.
- A final check ensures a response is always generated, or an explicit error is thrown for unresolvable queries, preventing silent failures.

To test the human escalation, send a query that won't be in our mock knowledge base and is likely classified as `complex_issue` or `unknown`:

```
npx trigger.dev@v4-beta send customer.query '{"userId": "user-456", "query":
"I need help configuring my new smart home device with the app, it keeps
failing. Can you help me troubleshoot?"}'
```

The workflow will pause at the `wait-for-human-input` step. Go to your Trigger.dev dashboard, find the running job, and you'll see it waiting. To resume it, you need to send an event. You can simulate this from your terminal:

```
npx trigger.dev@v4-beta send human.input '{"runId": "PASTE_YOUR_RUN_ID_HERE",
"response": "Hello! I understand your smart home device setup is challenging.
Please ensure your device is in pairing mode and your app is updated. If
issues persist, try restarting your router. Let me know if that helps!",
"approved": true}'
```

CRITICAL: Replace `PASTE_YOUR_RUN_ID_HERE` with the actual `run.id` from your Trigger.dev dashboard for the paused job. After sending this, the job will resume and complete, incorporating the human's response.

Step 4: Sending the Response to the Customer

The final step in our workflow is to deliver the generated or human-approved response back to the customer. This could be via email, a messaging platform, or a webhook to your frontend application.

Complete the `run` function in `src/jobs/customerSupportAgent.ts` by adding the final `io.runTask` call after the last `if (!responseToCustomer)` block:

```
// src/jobs/customerSupportAgent.ts (continued)
// ... (previous code for imports, client, OpenAI, ClassificationSchema,
fetchKnowledgeBase)

client.defineJob({
  id: "customer-support-agent",
  name: "AI Customer Support Agent",
  version: "1.0.0",
  enabled: true,
  trigger: client.defineEventTrigger({
    name: "customer.query",
    schema: z.object({
      userId: z.string(),
      query: z.string(),
      conversationId: z.string().optional(),
    }),
  }),
  run: async (payload, io, ctx) => {
    io.logger.info("Received customer query", payload);

    const classification = await io.runTask(
      "classify-query",
      {
        name: "Classify Customer Query with AI",
        icon: "openai",
        params: { query: payload.query },
      },
      async (task, params) => {
        const response = await openai.chat.completions.create({
          model: "gpt-4o",
          messages: [
            { role: "system", content: `You are a helpful assistant that
classifies customer queries into one of these categories: simple_faq,
complex_issue, billing_query, feature_request, or unknown. Provide relevant
keywords, a confidence score (0-1), and a brief summary. Respond strictly in
JSON format.` },
            { role: "user", content: `Classify the following query: "$
{params.query}"` },
          ],
          response_format: { type: "json_object" },
          temperature: 0.1,
        });

        const rawResult = response.choices[0].message?.content;
        if (!rawResult) {
          throw new Error("AI classification failed to return content.");
        }
        const parsedResult = JSON.parse(rawResult);
        const validatedClassification = ClassificationSchema.parse(parsedResul
t);
        return validatedClassification;
      }
    );

    io.logger.info("Query classified", { classification });

    let responseToCustomer: string | null = null;
    let humanProvidedResponse = false;

    if (classification.type === "simple_faq" && classification.confidence > 0.
7) {
```

```
      io.logger.info("Attempting to answer simple FAQ from knowledge base.");
      const kbAnswer = await io.runTask(
        "knowledge-base-lookup",
        {
          name: "Lookup Answer in Knowledge Base",
          icon: "database",
          params: { query: payload.query },
        },
        async (task, params) => {
          return await fetchKnowledgeBase(params.query);
        }
      );

      if (kbAnswer) {
        responseToCustomer = kbAnswer;
        io.logger.info("Found answer in KB.", { answer: kbAnswer });
      } else {
        io.logger.warn("Could not find answer in KB, will consider human
escalation.");
      }
    }

    if (!responseToCustomer) {
      io.logger.info("Escalating to human agent for review.");

      await io.runTask(
        "notify-human-agent",
        {
          name: "Notify Human Agent of Escalation",
          icon: "bell",
          params: {
            userId: payload.userId,
            query: payload.query,
            classification: classification.summary || classification.type,
            runId: ctx.run.id,
          },
        },
        async (task, params) => {
          io.logger.warn(`Human Agent NOTIFIED: User ${params.userId} needs
help with: ${params.query}. Run ID: ${params.runId}`);
          return `https://your-agent-dashboard.com/review/${ctx.run.id}`;
        }
      );

      const humanInput = await io.wait(
        "wait-for-human-input",
        {
          name: "Wait for Human Agent Input",
          timeoutInSeconds: 60 * 60 * 24 * 7,
          schema: z.object({
            response: z.string().min(1, "Response cannot be empty."),
            approved: z.boolean().optional().default(true),
          }),
        }
      );

      io.logger.info("Human agent provided input.", { humanInput });
      if (humanInput.approved) {
        responseToCustomer = humanInput.response;
        humanProvidedResponse = true;
      } else {
        responseToCustomer = "The human agent reviewed your query. " + humanIn
```

```
put.response;
        io.logger.warn("Human agent disapproved the automated response, but
provided an alternative.");
      }
    }

    if (!responseToCustomer) {
      io.logger.error("Failed to generate a response, even after human
escalation.");
      throw new Error("Unable to resolve customer query automatically or with
human input.");
    }

    // 🔥 Optimization / Pro tip: Use a dedicated messaging service or webhook
here.
    // This task represents the final delivery of the response to the
customer.
    await io.runTask(
      "send-response-to-customer",
      {
        name: "Send Response to Customer",
        icon: "send", // Use a send icon
        params: {
          userId: payload.userId,
          message: responseToCustomer,
          source: humanProvidedResponse ? "human" :
"ai", // Indicate if human or AI generated
          conversationId: payload.conversationId,
        },
      },
      async (task, params) => {
        // In a real application, this would send an email (e.g., via
SendGrid),
        // push a message to a chat UI (e.g., via WebSockets), or call a
webhook
        // to update a CRM or a custom frontend.
        io.logger.info(`Sending response to user ${params.userId}: ${params.me
ssage}`);
        // Simulate a successful send operation
        return { success: true, messageId: `msg-${Date.now()}` };
      }
    );

    io.logger.info("Workflow completed successfully!");
    // Return a final status and the response for observability
    return { status: "resolved", finalResponse: responseToCustomer };
  },
});
```

Explanation:

- The final `io.runTask("send-response-to-customer", ...)` simulates delivering the message back to the customer. This would likely involve integrating with a specific messaging API (e.g., Twilio for SMS, SendGrid for email, or a custom WebSocket service for real-time chat in a web application).

- The `source` parameter helps track whether the response originated from the AI directly or was approved/modified by a human, which can be valuable for analytics and auditing.
- The workflow concludes by logging its success and returning a final status and the response, making the outcome clear in the Trigger.dev dashboard.

Now, you have a complete AI-powered customer support agent workflow! Test it with various queries to observe the different paths: direct AI response from the knowledge base, or human escalation followed by a human-approved response.

Mini-Challenge: Contextual Conversations

Our current agent handles each customer query in isolation. However, real customer support conversations are contextual, building upon previous turns.

Challenge: Modify the `customer-support-agent` workflow to maintain a simple conversation history. When a new query comes in, if it's part of an existing `conversationId`, retrieve the previous turn(s) and include them in the AI classification prompt to give the AI more context.

Hint:

- You'll need `io.store.set` to save the current query and response (and optionally the classification) at the end of a job run, associated with the `conversationId`.
- You'll need `io.store.get` to retrieve past conversation turns associated with a `conversationId` at the beginning of a new job run.
- The `payload` already includes an optional `conversationId`.
- Pass the retrieved history as part of the `messages` array to the OpenAI API when classifying the new query. Remember to limit the history to a reasonable number of turns (e.g., the last 3-5 turns) to stay within token limits and manage costs. Consider structuring the history as `[{ role: "user", content: "..." }, { role: "assistant", content: "..." }]`.

What to observe/learn: This challenge will teach you how to use Trigger.dev's durable storage (`io.store`) to maintain state across workflow runs, enabling long-running, stateful interactions like multi-turn conversations. It also reinforces the importance of prompt engineering for contextual AI.

Common Pitfalls & Troubleshooting

Building distributed AI workflows, especially those involving human interaction, can introduce new challenges. Here are some common pitfalls and how to troubleshoot them with Trigger.dev.

1. **State Management in Long-Running Workflows:**
    - **Pitfall:** Forgetting to persist critical data (like conversation history, user preferences, or intermediate results) using Trigger.dev's `io.store` or passing it explicitly between `io.runTask` calls. If your workflow involves `io.wait` or `io.sleep`, any in-memory state will be lost when the workflow pauses and resumes.
    - **Troubleshooting:** Always assume that anything not explicitly stored or passed will be lost across `io.wait` boundaries or between distinct job runs.
        - Use `io.store.set` and `io.store.get` for state that needs to persist across pauses or multiple workflow runs (e.g., conversation history associated with a `conversationId`).
        - For transient data within a single `io.runTask` or between synchronous steps within the same `run` function execution, local variables are perfectly fine.
2. **API Rate Limits with AI Services:**
    - **Pitfall:** Hitting rate limits on external AI APIs (like OpenAI) during peak usage or when processing many concurrent requests, leading to `429 Too Many Requests` errors. This can interrupt your workflow.
    - **Troubleshooting:** Trigger.dev's `io.runTask` automatically handles retries with exponential backoff for transient errors, which is often sufficient for occasional rate limit issues. For very high throughput or sustained rate limiting, consider:
        - Implementing client-side rate limiting before sending events to Trigger.dev, if your system generates events very rapidly.
        - Using batch processing if the AI API supports it (e.g., sending multiple classification requests in one API call).
        - Distributing requests across multiple AI API keys or accounts if your provider and license allow it.
        - Optimizing AI prompts to reduce token usage and improve response times, which can indirectly help with rate limits.

3. **Debugging Asynchronous and Distributed Workflows:**
    - **Pitfall:** It can be hard to trace the flow of execution, especially when `io.wait` pauses a workflow for an extended period, or when issues occur across multiple `io.runTask` calls and external services. Traditional `console.log` debugging is insufficient.
    - **Troubleshooting:**
        - **Trigger.dev Dashboard:** This is your primary and most powerful tool. Every `io.runTask`, `io.wait`, and `io.logger` call is visible on a timeline, providing a clear, step-by-step trace of execution, including inputs, outputs, and any errors.
        - **`io.logger`:** Use `io.logger.info`, `warn`, `error` liberally to understand the state of your workflow at critical junctures. These logs appear directly in the dashboard, correlated with the specific task and run.
        - **`ctx.run.id`:** Use the `run.id` (available in the `ctx` object) to correlate logs and events across different systems if you're integrating with external monitoring tools or CRMs. Pass this ID to external services so you can easily link their logs back to a specific Trigger.dev workflow run.

Summary

Congratulations! You've successfully built a sophisticated AI-powered customer support agent using Trigger.dev, integrating AI classification, knowledge base lookup, and human-in-the-loop escalation.

Here are the key takeaways from this project:

- **AI Agent Orchestration:** Trigger.dev provides a robust and observable platform for orchestrating complex AI agent workflows, seamlessly integrating AI inference with business logic and external services.
- **Human-in-the-Loop (HITL):** You learned how to implement essential HITL patterns using `io.wait`, ensuring human oversight for critical decisions and ambiguous cases, building trust and improving accuracy.
- **Durable Execution:** The agent leveraged Trigger.dev's durable execution to maintain state and context across long-running, multi-step processes, including pauses for human input, without losing progress.

- **Robustness:** Features like automatic retries (via `io.runTask`) for transient failures and clear observability in the Trigger.dev dashboard contribute to a highly resilient and reliable system.
- **Modular Design:** By breaking down the agent's logic into distinct `io.runTask` calls and conditional paths, we created a modular, easy-to-understand, and maintainable workflow.
- **Schema Validation:** Using Zod for validating AI outputs is a critical best practice for building robust AI-driven applications, ensuring predictable data flow.

This project demonstrates how Trigger.dev empowers developers to build intelligent, reliable, and scalable automated systems that combine the power of AI with necessary human intervention. From here, you can further enhance this agent with more complex AI reasoning, integrations with real CRMs, and rich user interfaces for both customers and human agents.

References

- Trigger.dev Documentation
- Trigger.dev GitHub Repository
- OpenAI API Documentation
- Zod Documentation

This page is AI-assisted and reviewed. It references official documentation and recognized resources where relevant.