

# Hashnode's journey into its public GraphQL API

From Next.js API routes to a serverless GraphQL API



Author

SANDRO VOLPICELLA

Date

2024-10-03

# whoami


- Software Engineer (Fullstack Serverless) from Munich, Germany 🇩🇪🥨
- Platform Engineering Lead @hashnode since 2022
- AWS Community Builder 🧡
- 📖 [awsfundamentals.com](https://awsfundamentals.com) | 📖 [cloudwatchbook.com](https://cloudwatchbook.com)
- Plays tennis 🎾 & lifts wights 🏋️



Author

SANDRO VOLPICELLA

@SANDRO\_VOL





# What is Hashnode?

API & DEVELOPER DOCS | DEVELOPER BLOGGING PLATFORM | HEADLESS CMS

# Developer Blogging

- Superb editor + AI
- Blogging on custom domain
- Headless CMS → Bring your own frontend
- Newsletters, Webhooks, GitHub Backup, RSS, and more!

## Defining APIs with Lambda as Compute

Once your SST project is set up, the next step is to define our APIs. SST makes it **very easy** to create and deploy AWS Lambda functions.

But how does everything play together at SST? Let's go one step back.

It all starts in the `sst.config.ts` file:

```
TypeScript
import { SSTConfig } from 'sst';
import { Stack } from './stacks/index';

export default {
  config(_input) {
    return {
      name: 'bedrock-openai-experiments-chat',
      region: 'us-east-1',
    };
  },
  stacks(app) {
```

Try Pitch

# Developer Documentation

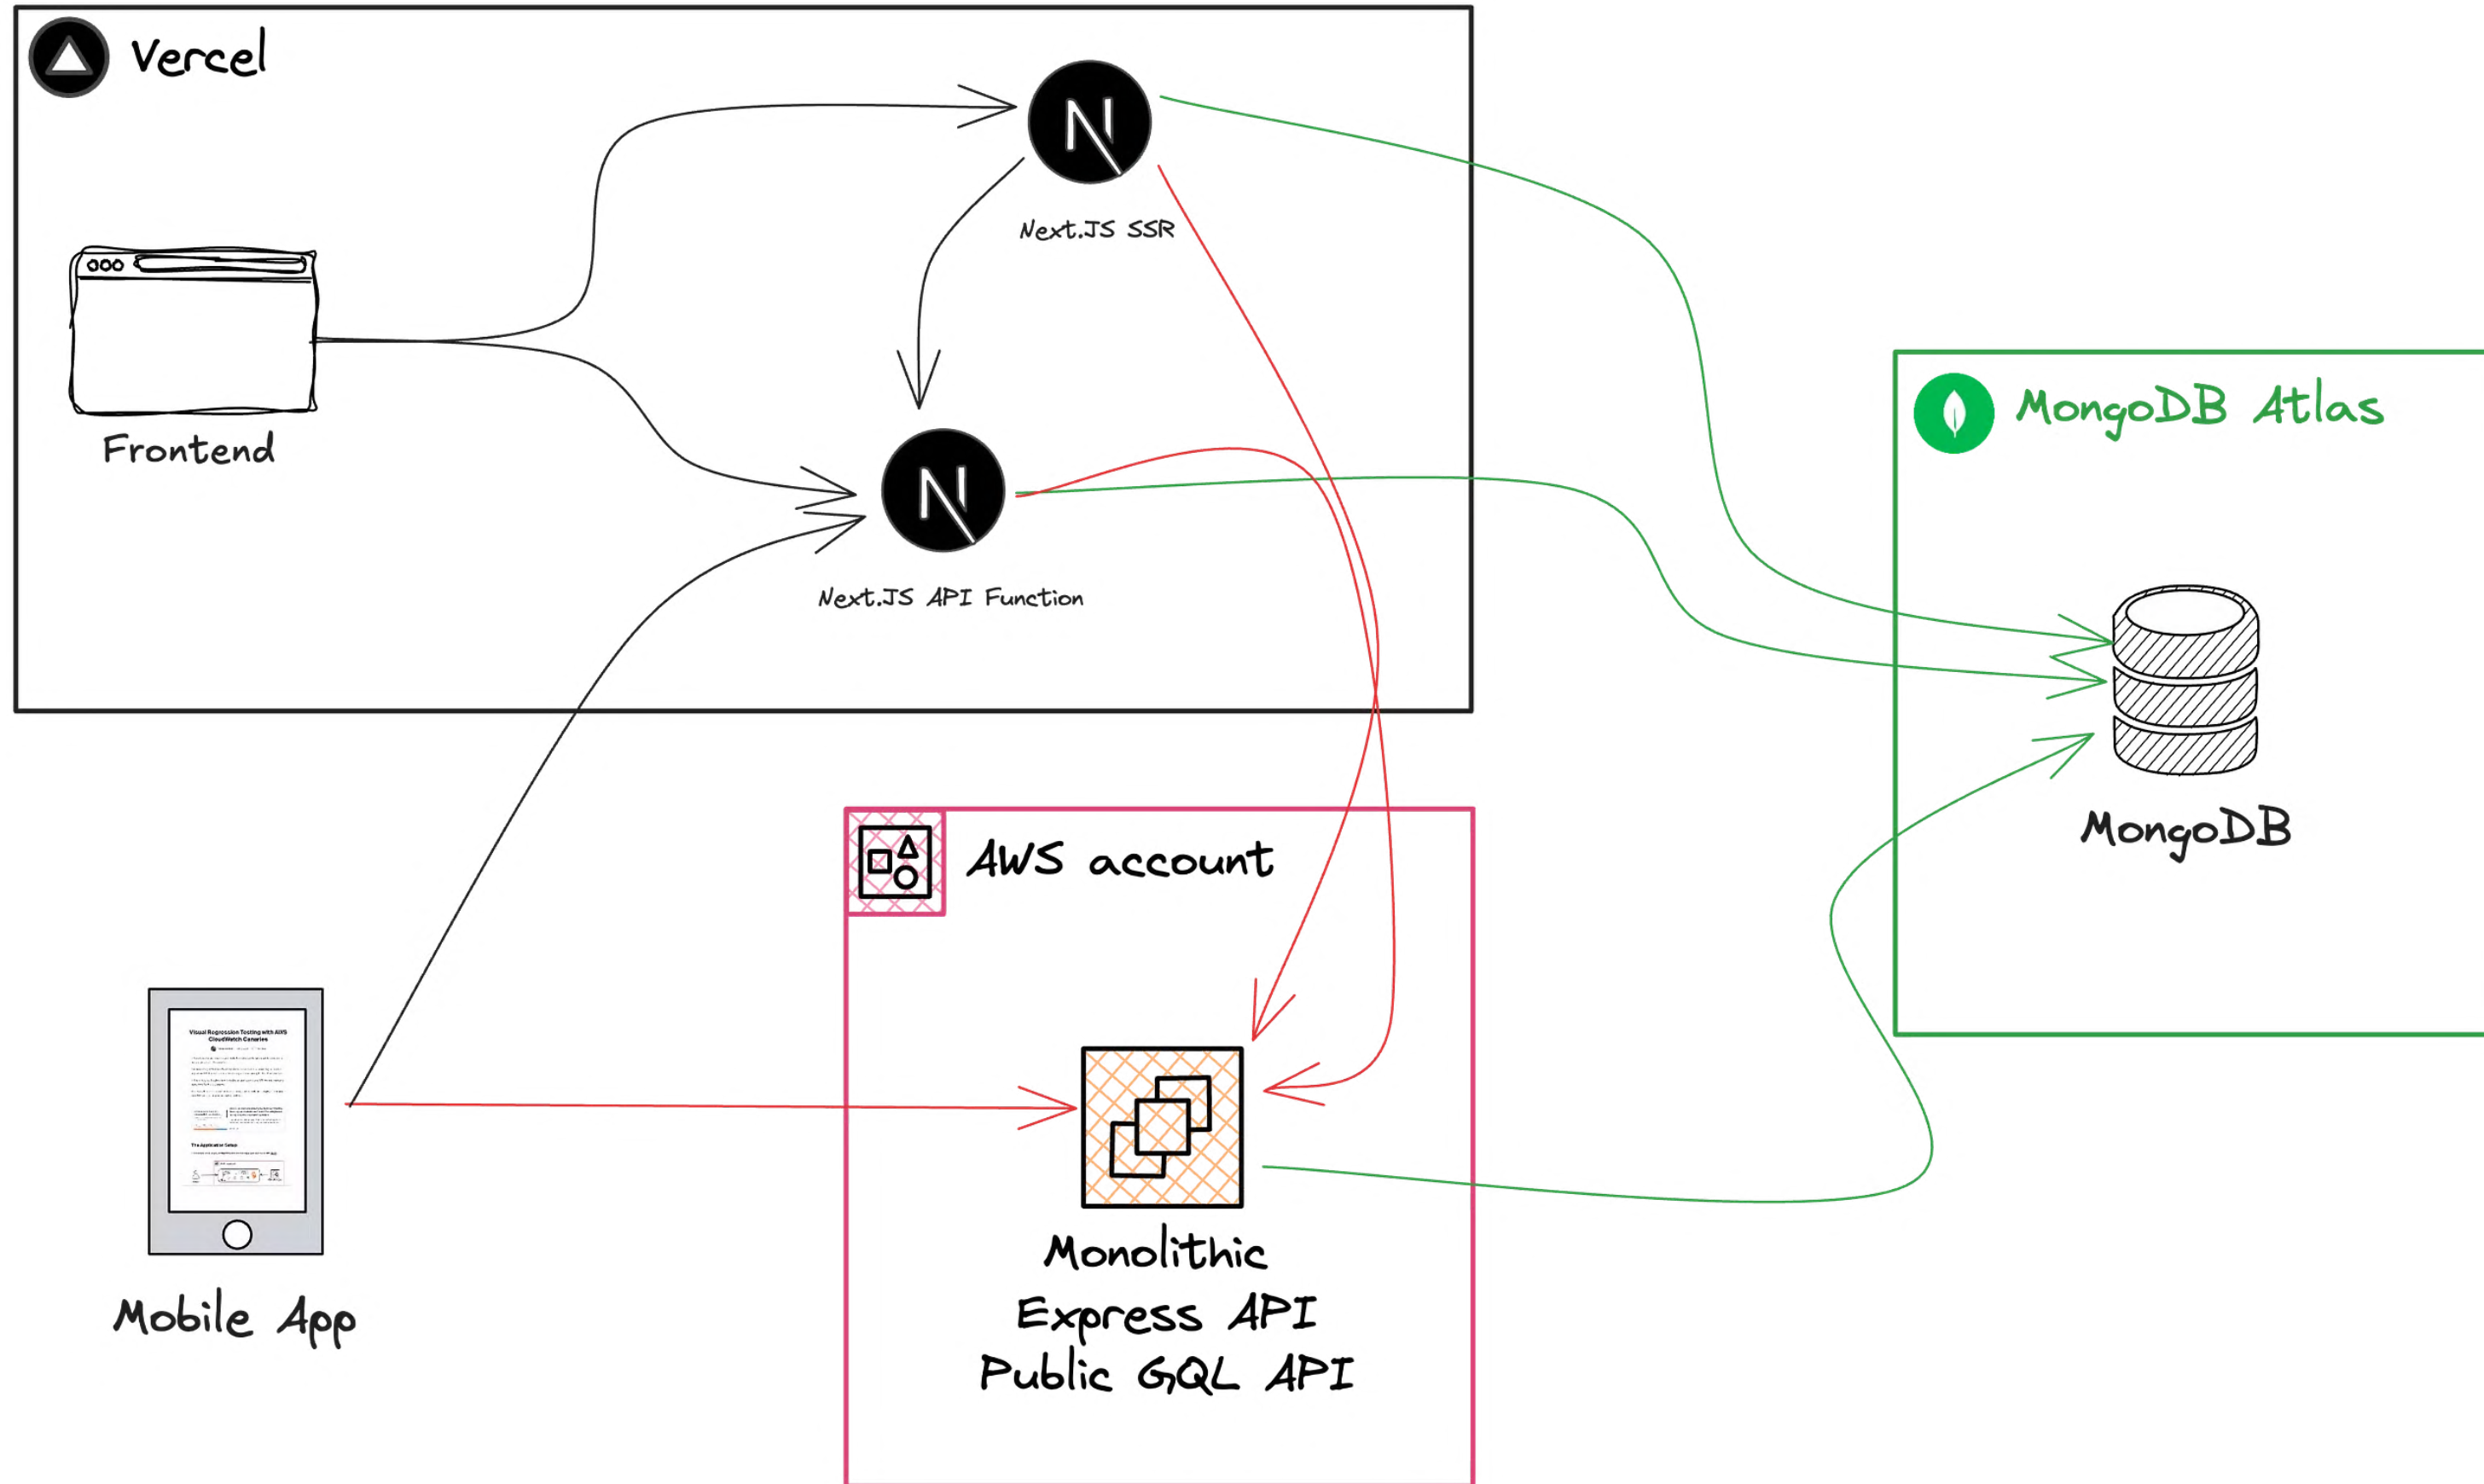
- MDX-powered Editor
- Open API/Swagger automatic API documentation
- Endless flexibility with headless CMS + GQL
- Fulltext + AI Search


The screenshot shows the Hashnode Public APIs GQL Playground interface. On the left is a sidebar with navigation links: Overview, GQL Playground (selected), Caching, Rate Limit, Authentication, Status and Error Codes, Pagination, Breaking Change, Legacy API Migration, Count, Example, Queries, documentationProject, checkSubdomainAvailability, checkCustomDomainAvailability, Feed, Draft, Me, Post, Publication, Scheduled Post, Search Posts of Publication, and Tag. The main content area is titled "GQL Playground" and contains the following text: "All Hashnode Public API queries are made through a single GraphQL endpoint, which only accepts POST requests." Below this is a link to "https://gql.hashnode.com". Further down, it says "You can visit the same URL to check out Hashnode API Playground." At the bottom, there is a preview of the GraphQL playground interface showing a query and its JSON response. The response includes fields like "name", "region", "description", "tags", "draft", "published", "scheduled", "searchable", "feed", "drafts", "posts", "tags", "postsCount", "draftsCount", "scheduledCount", "searchableCount", "feedCount", "draftsCount", "postsCount", "tagsCount", "postsCount", "draftsCount", "scheduledCount", "searchableCount", "feedCount", "draftsCount", "postsCount", "tagsCount".



# Starting with an API mess


MONOLITHIC EXPRESS APP | NEXT.JS API ROUTES ON VERCEL








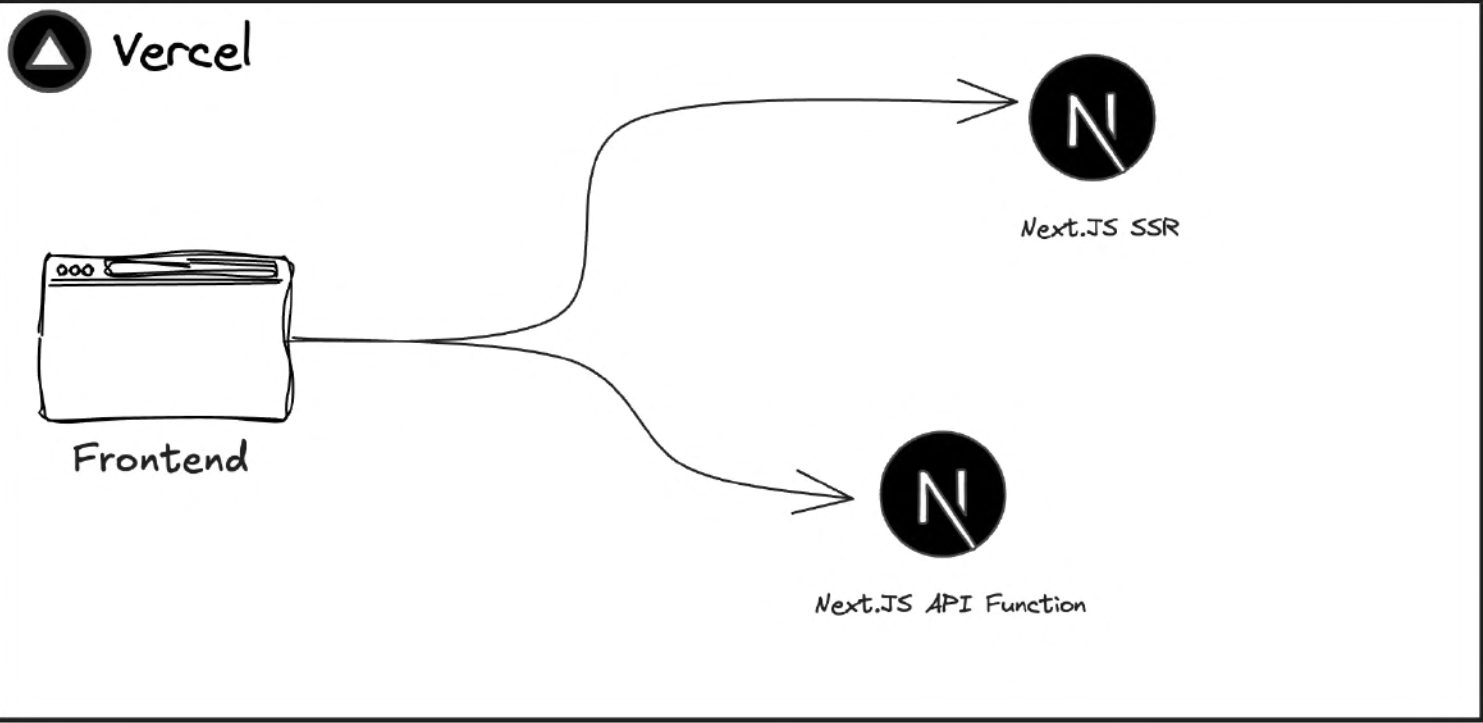
Vercel

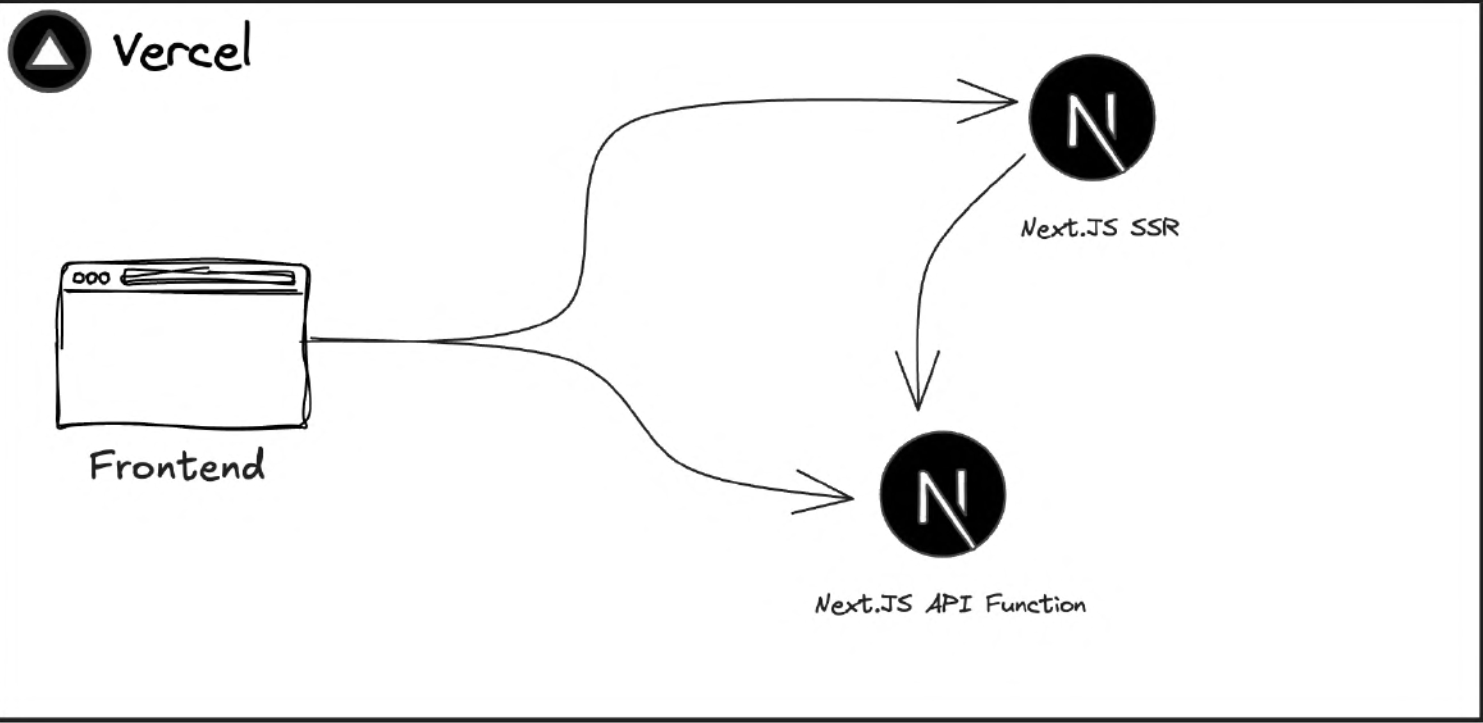



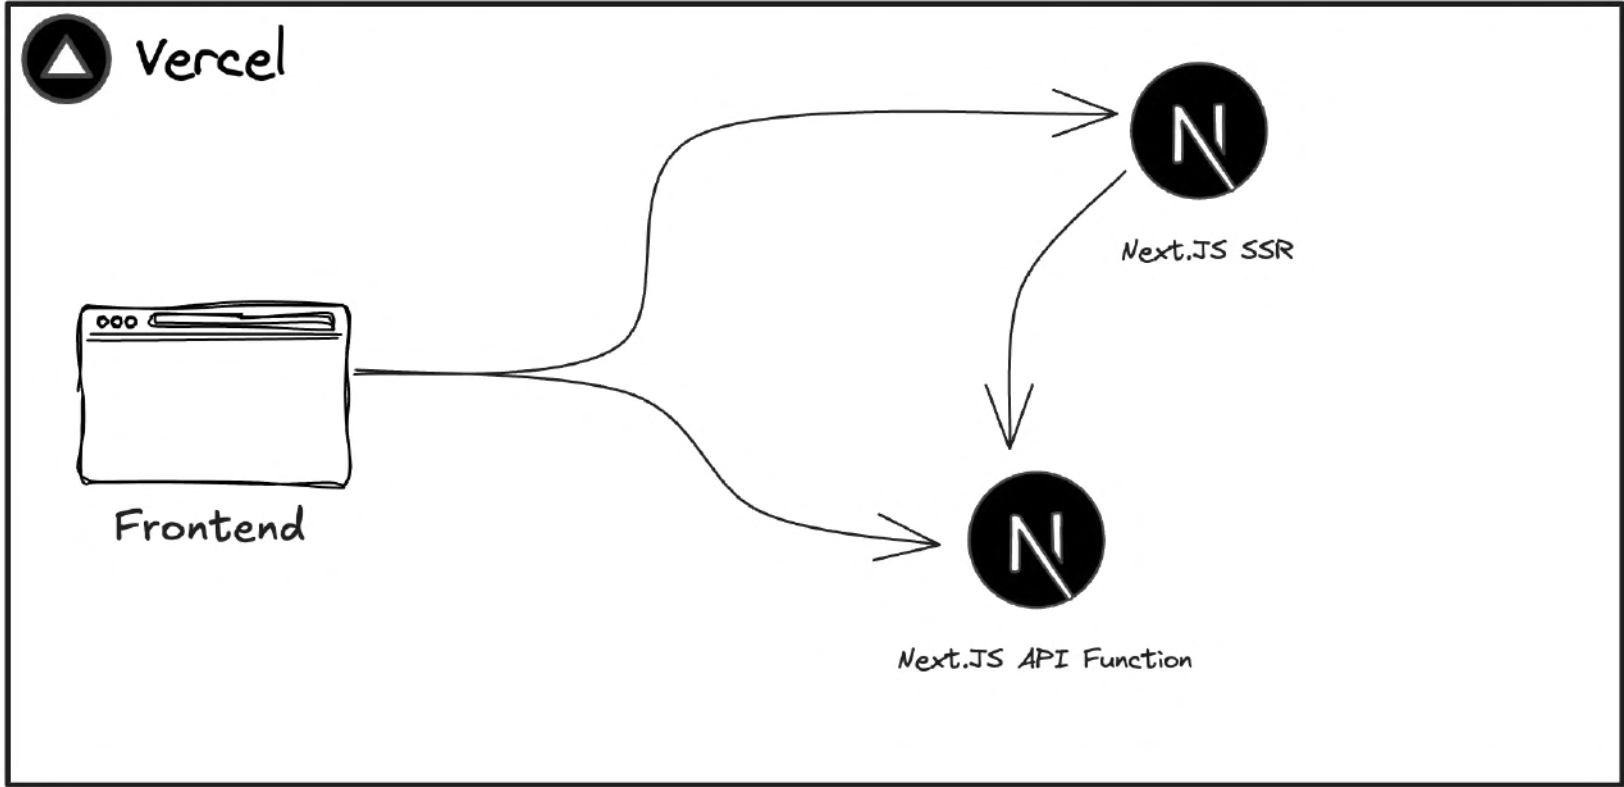
Frontend

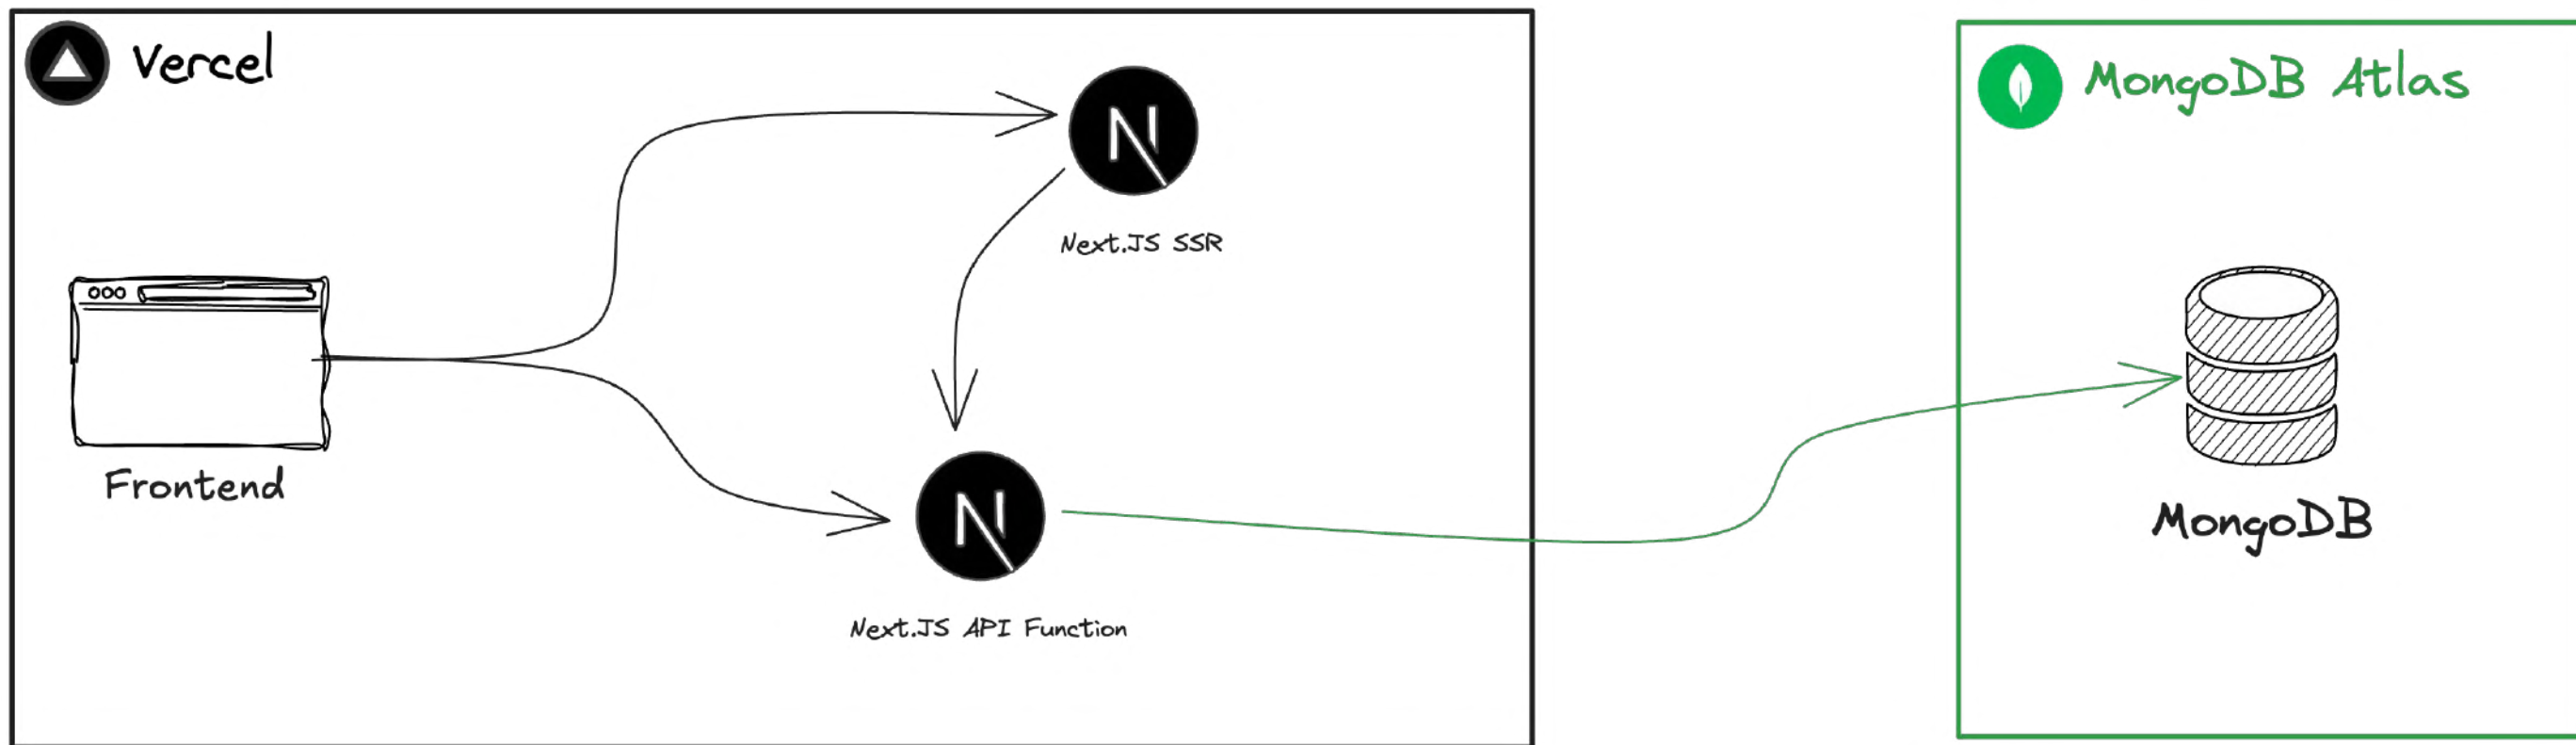


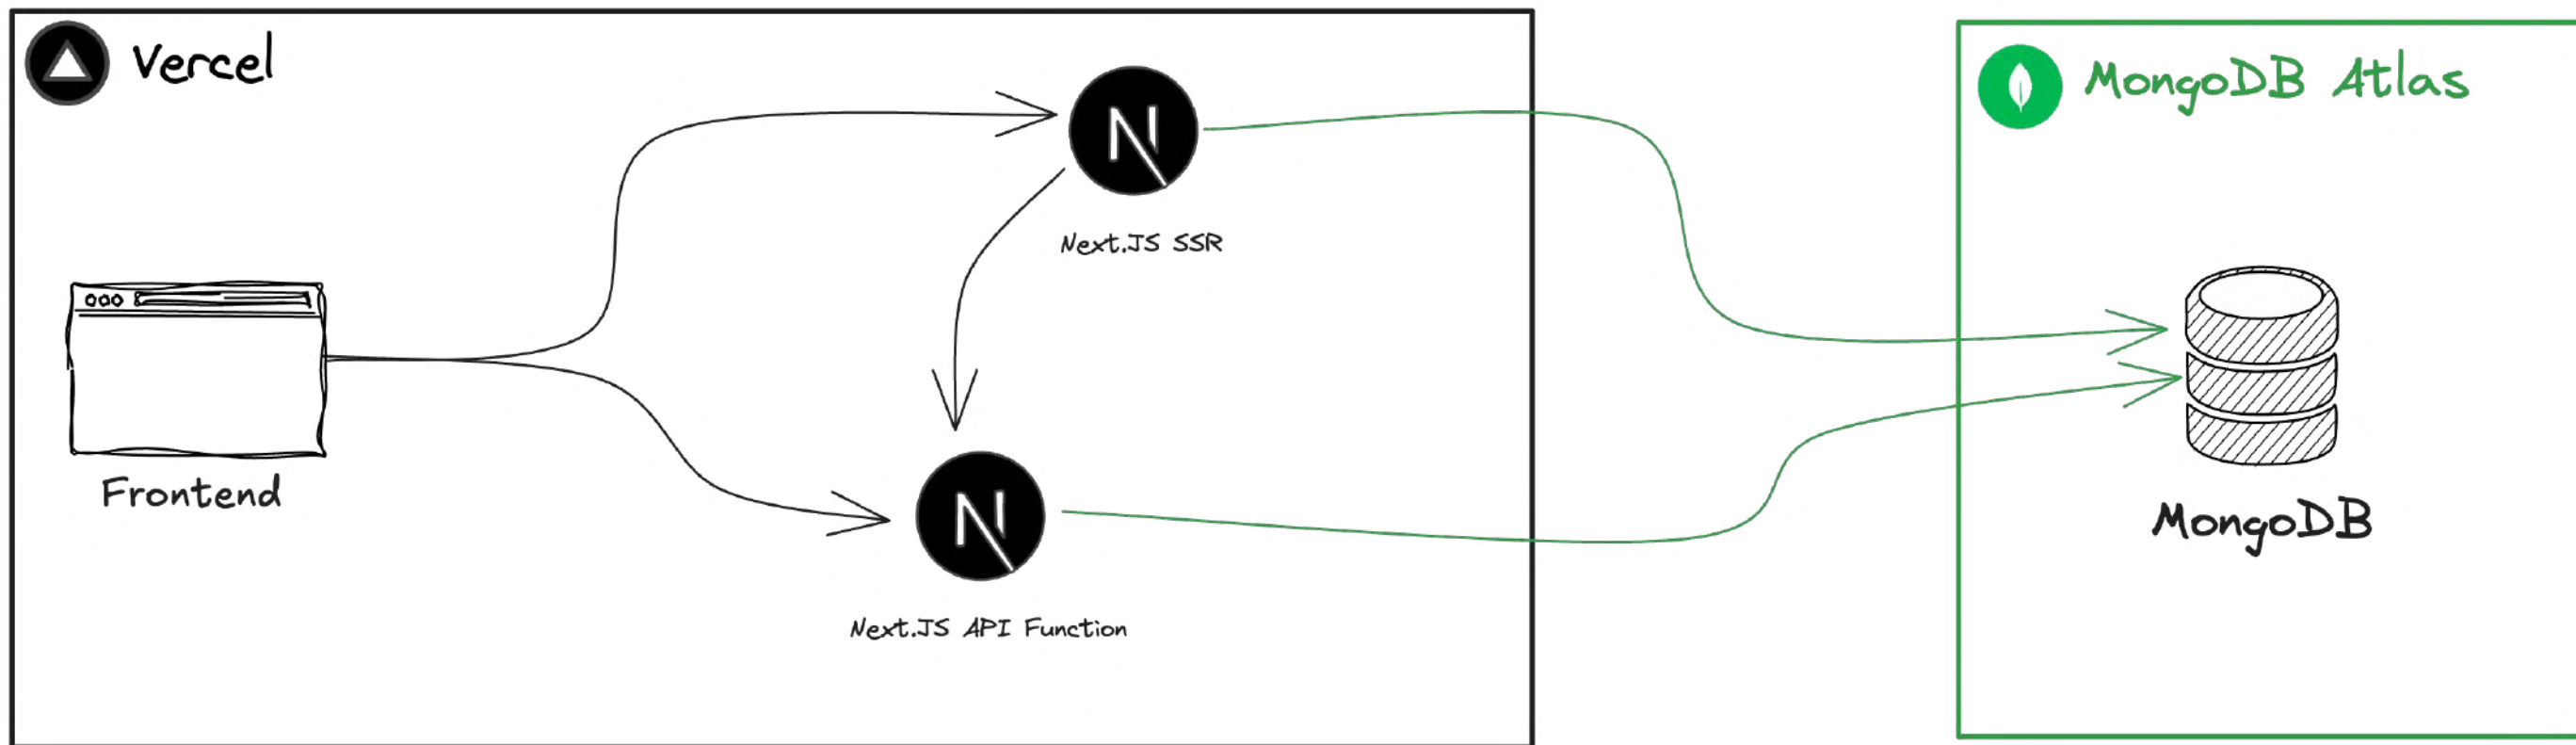


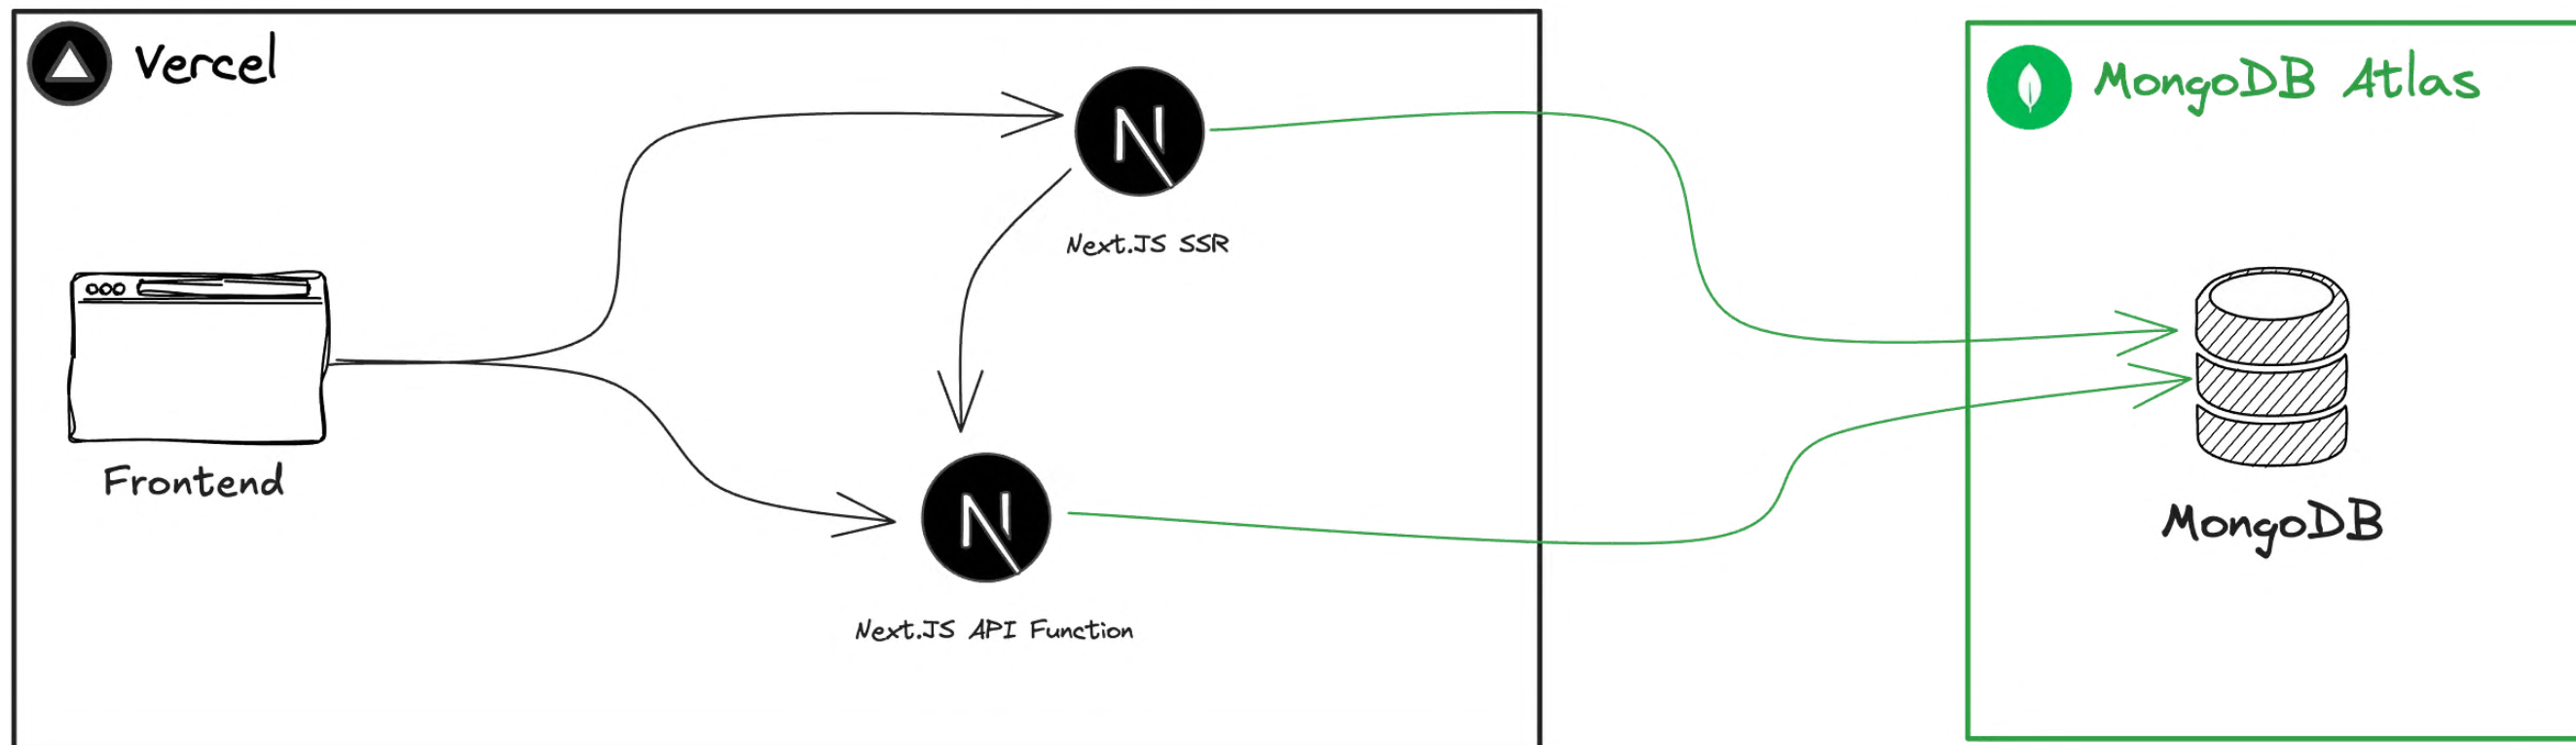



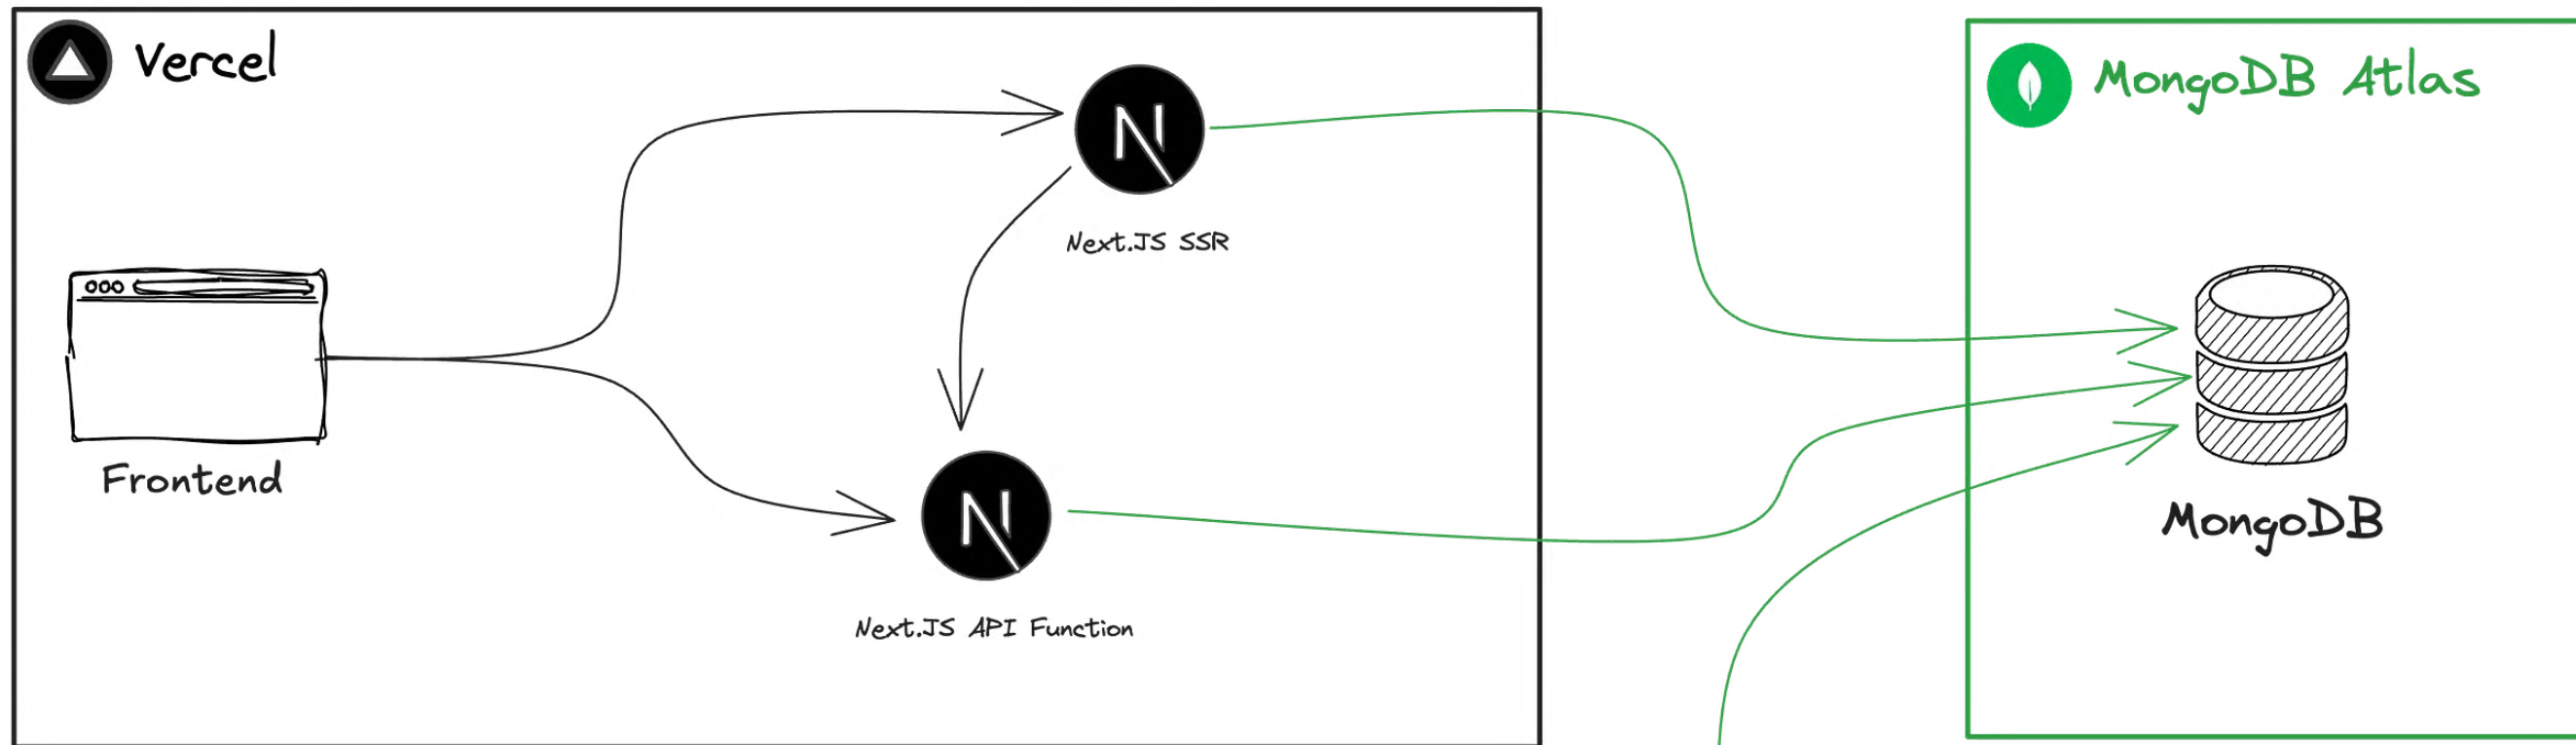




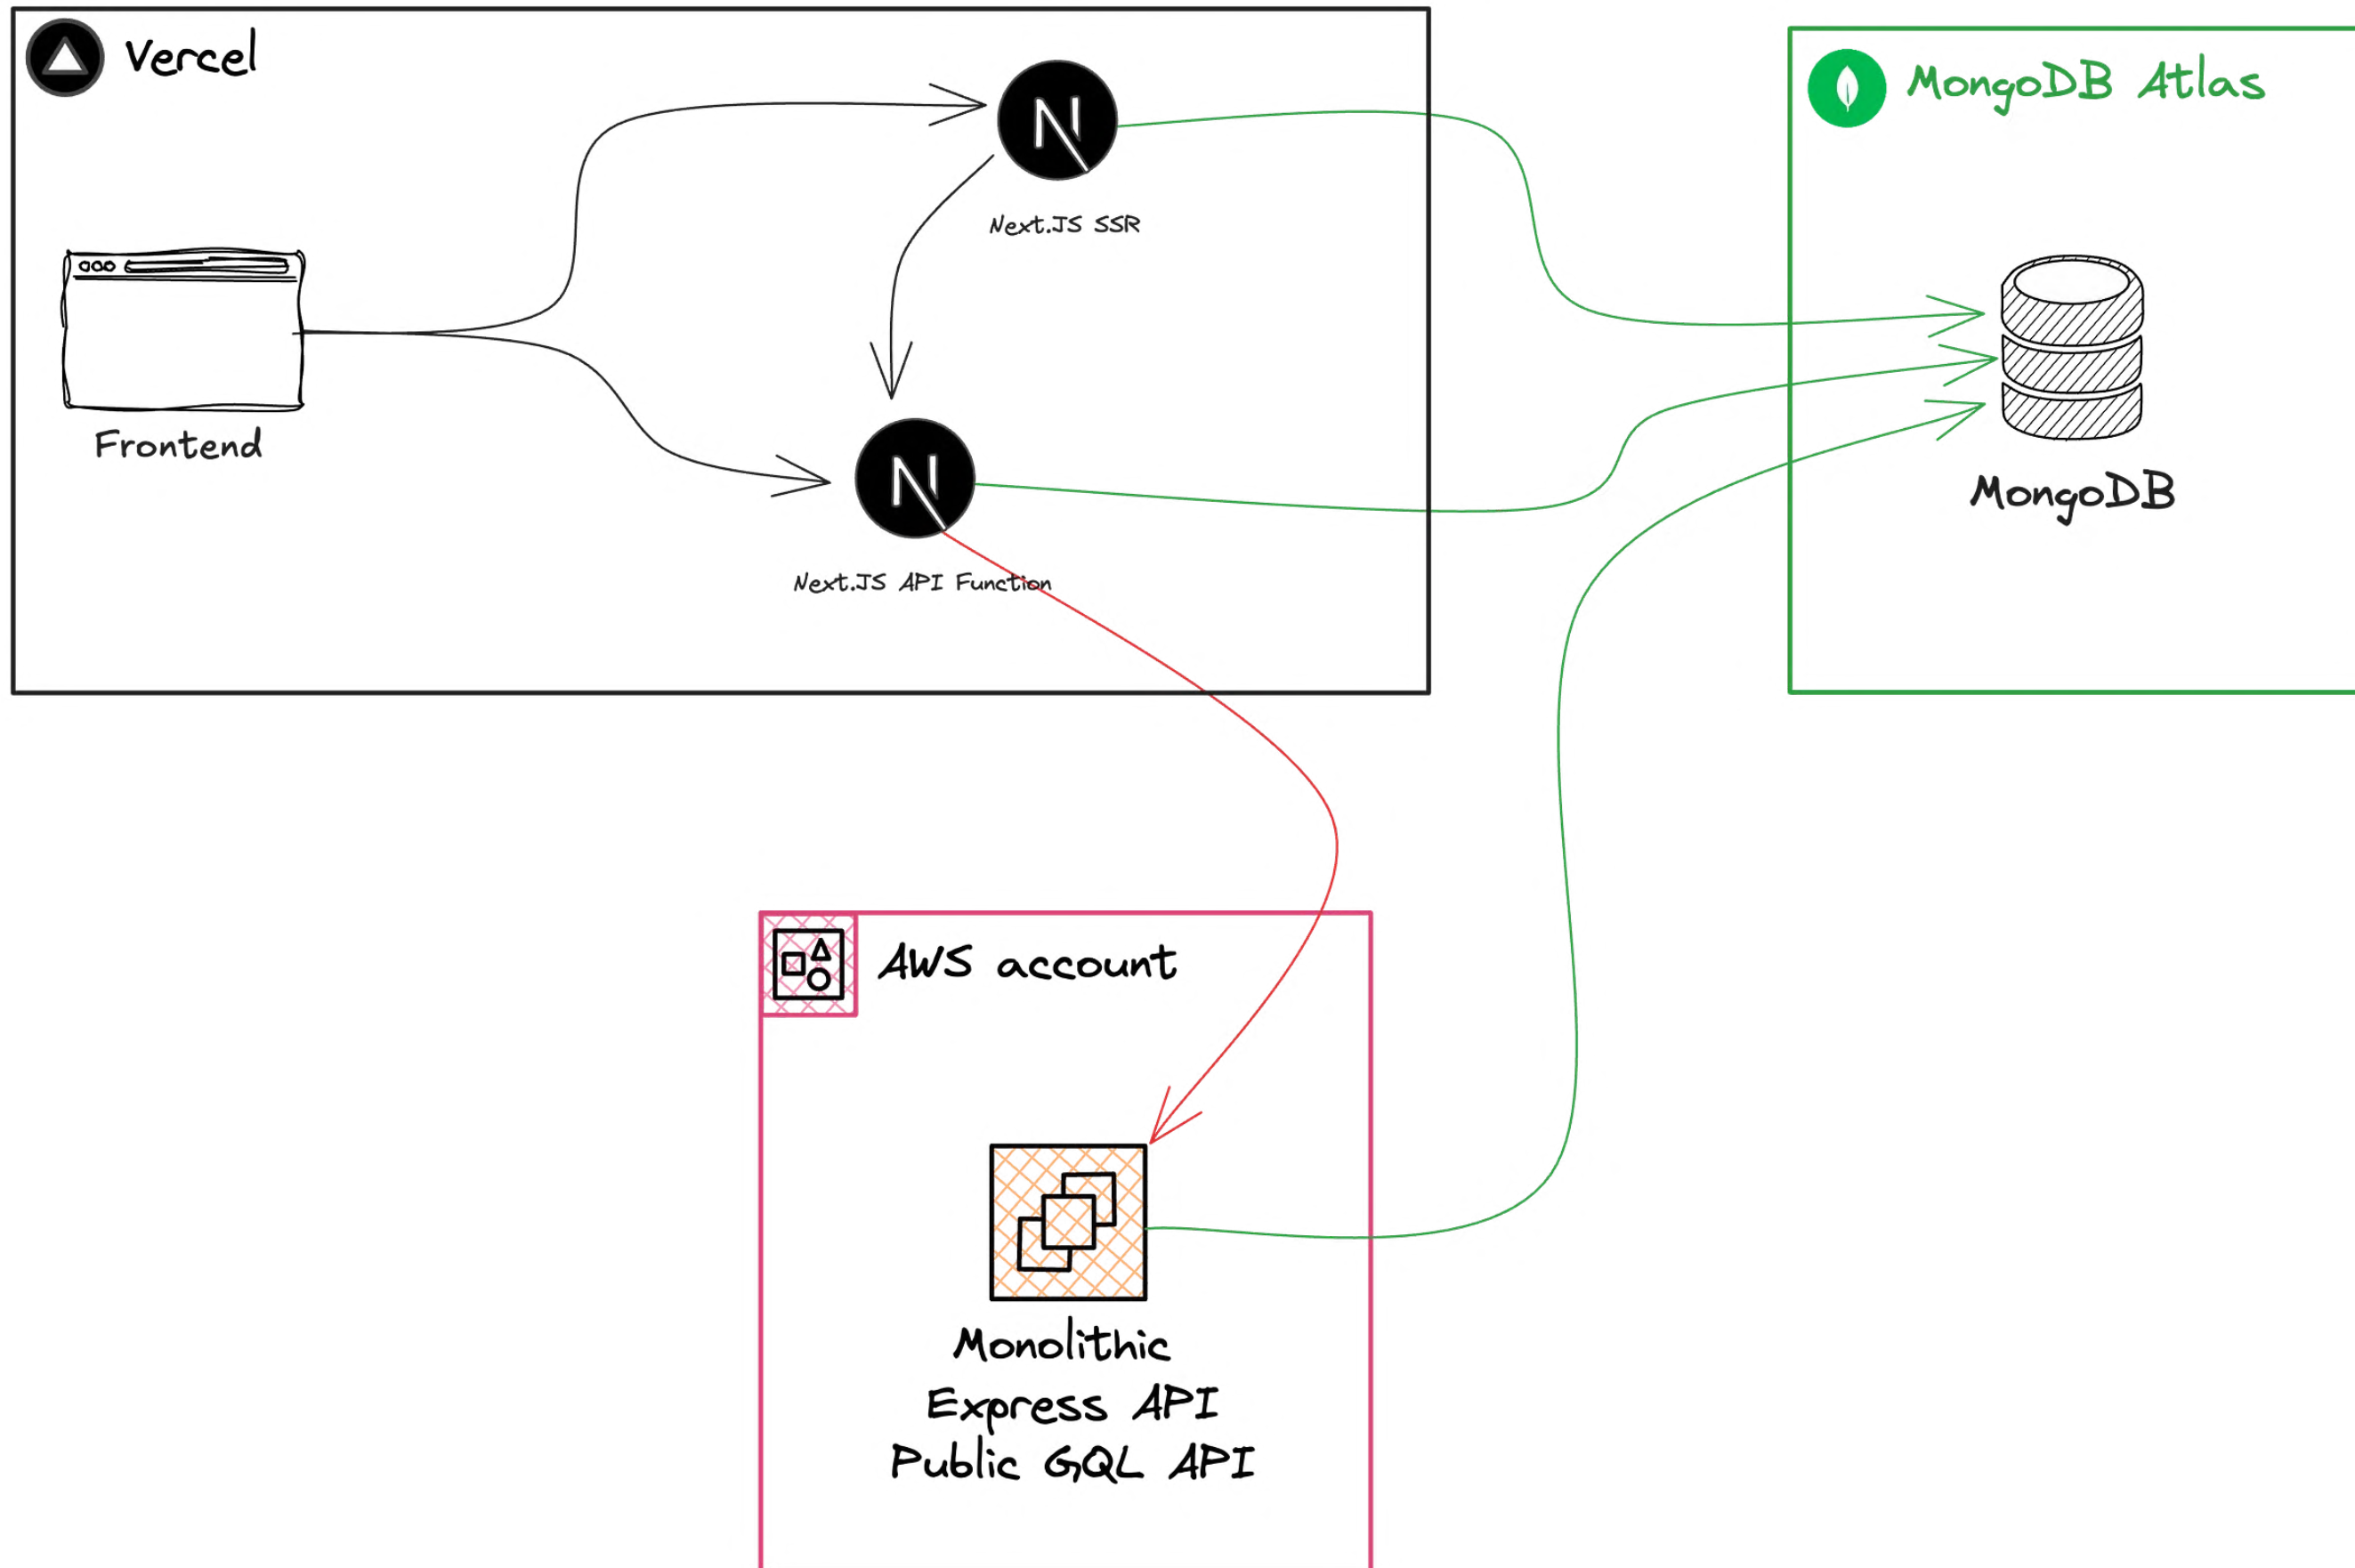


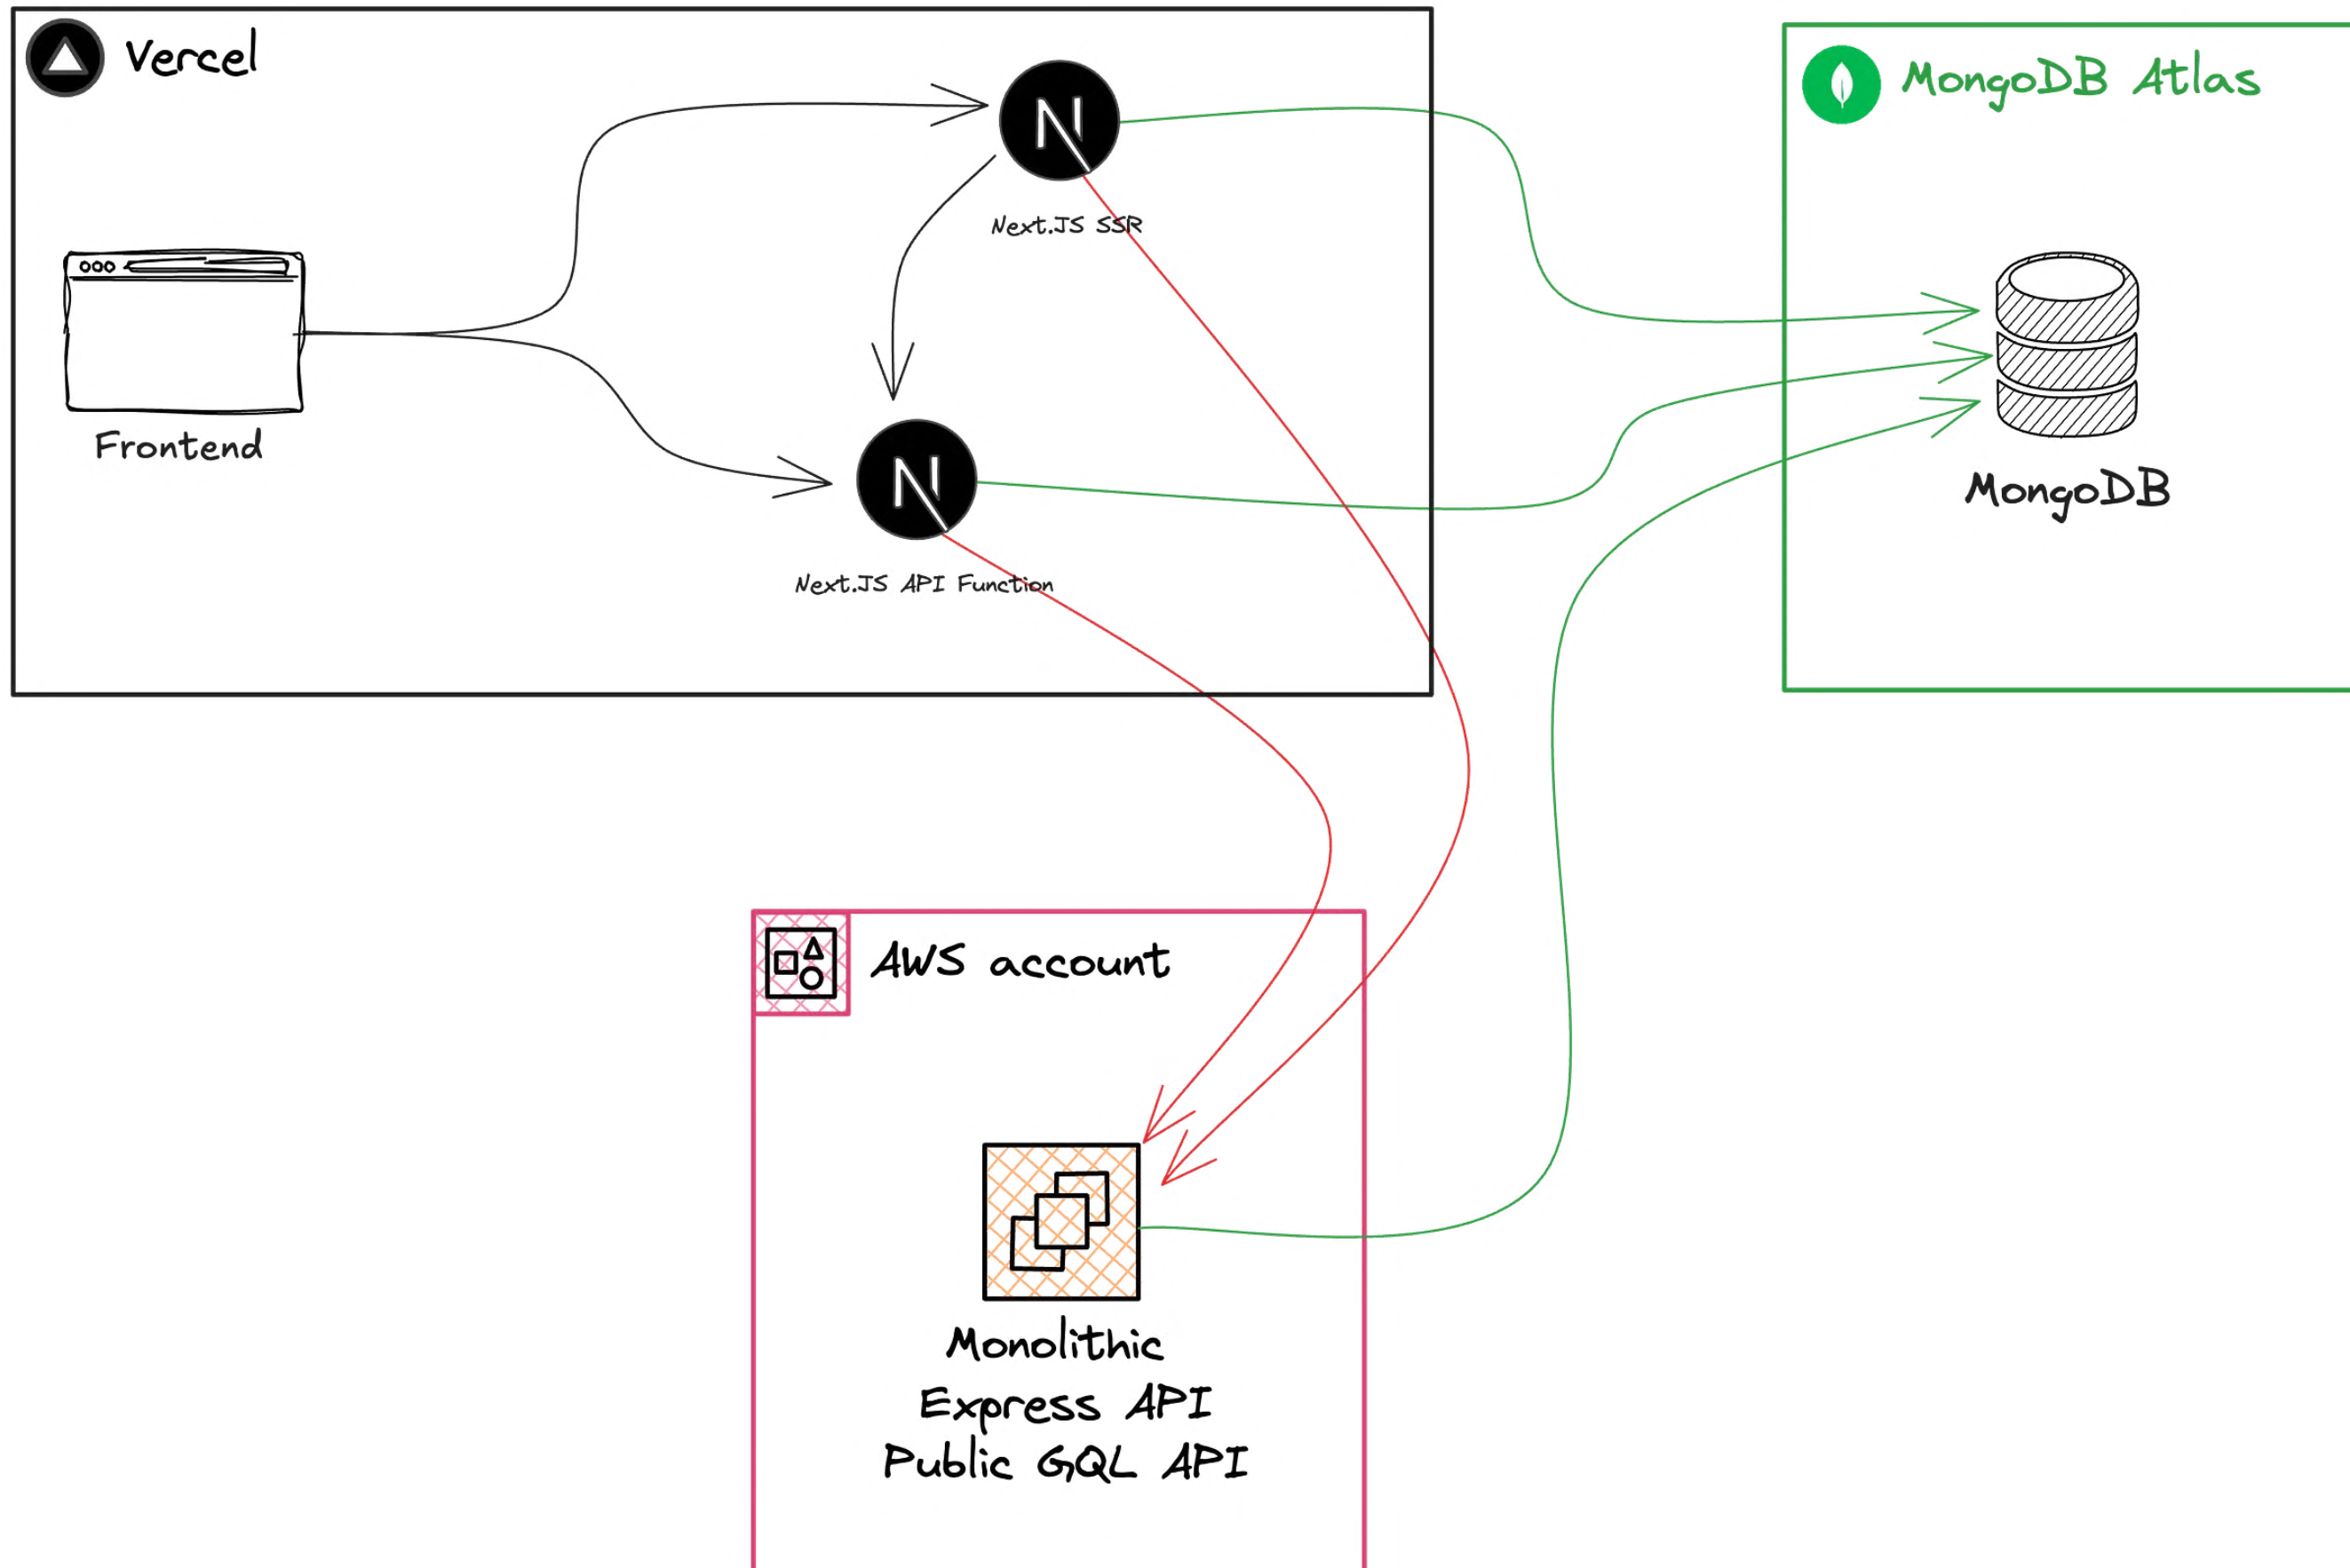


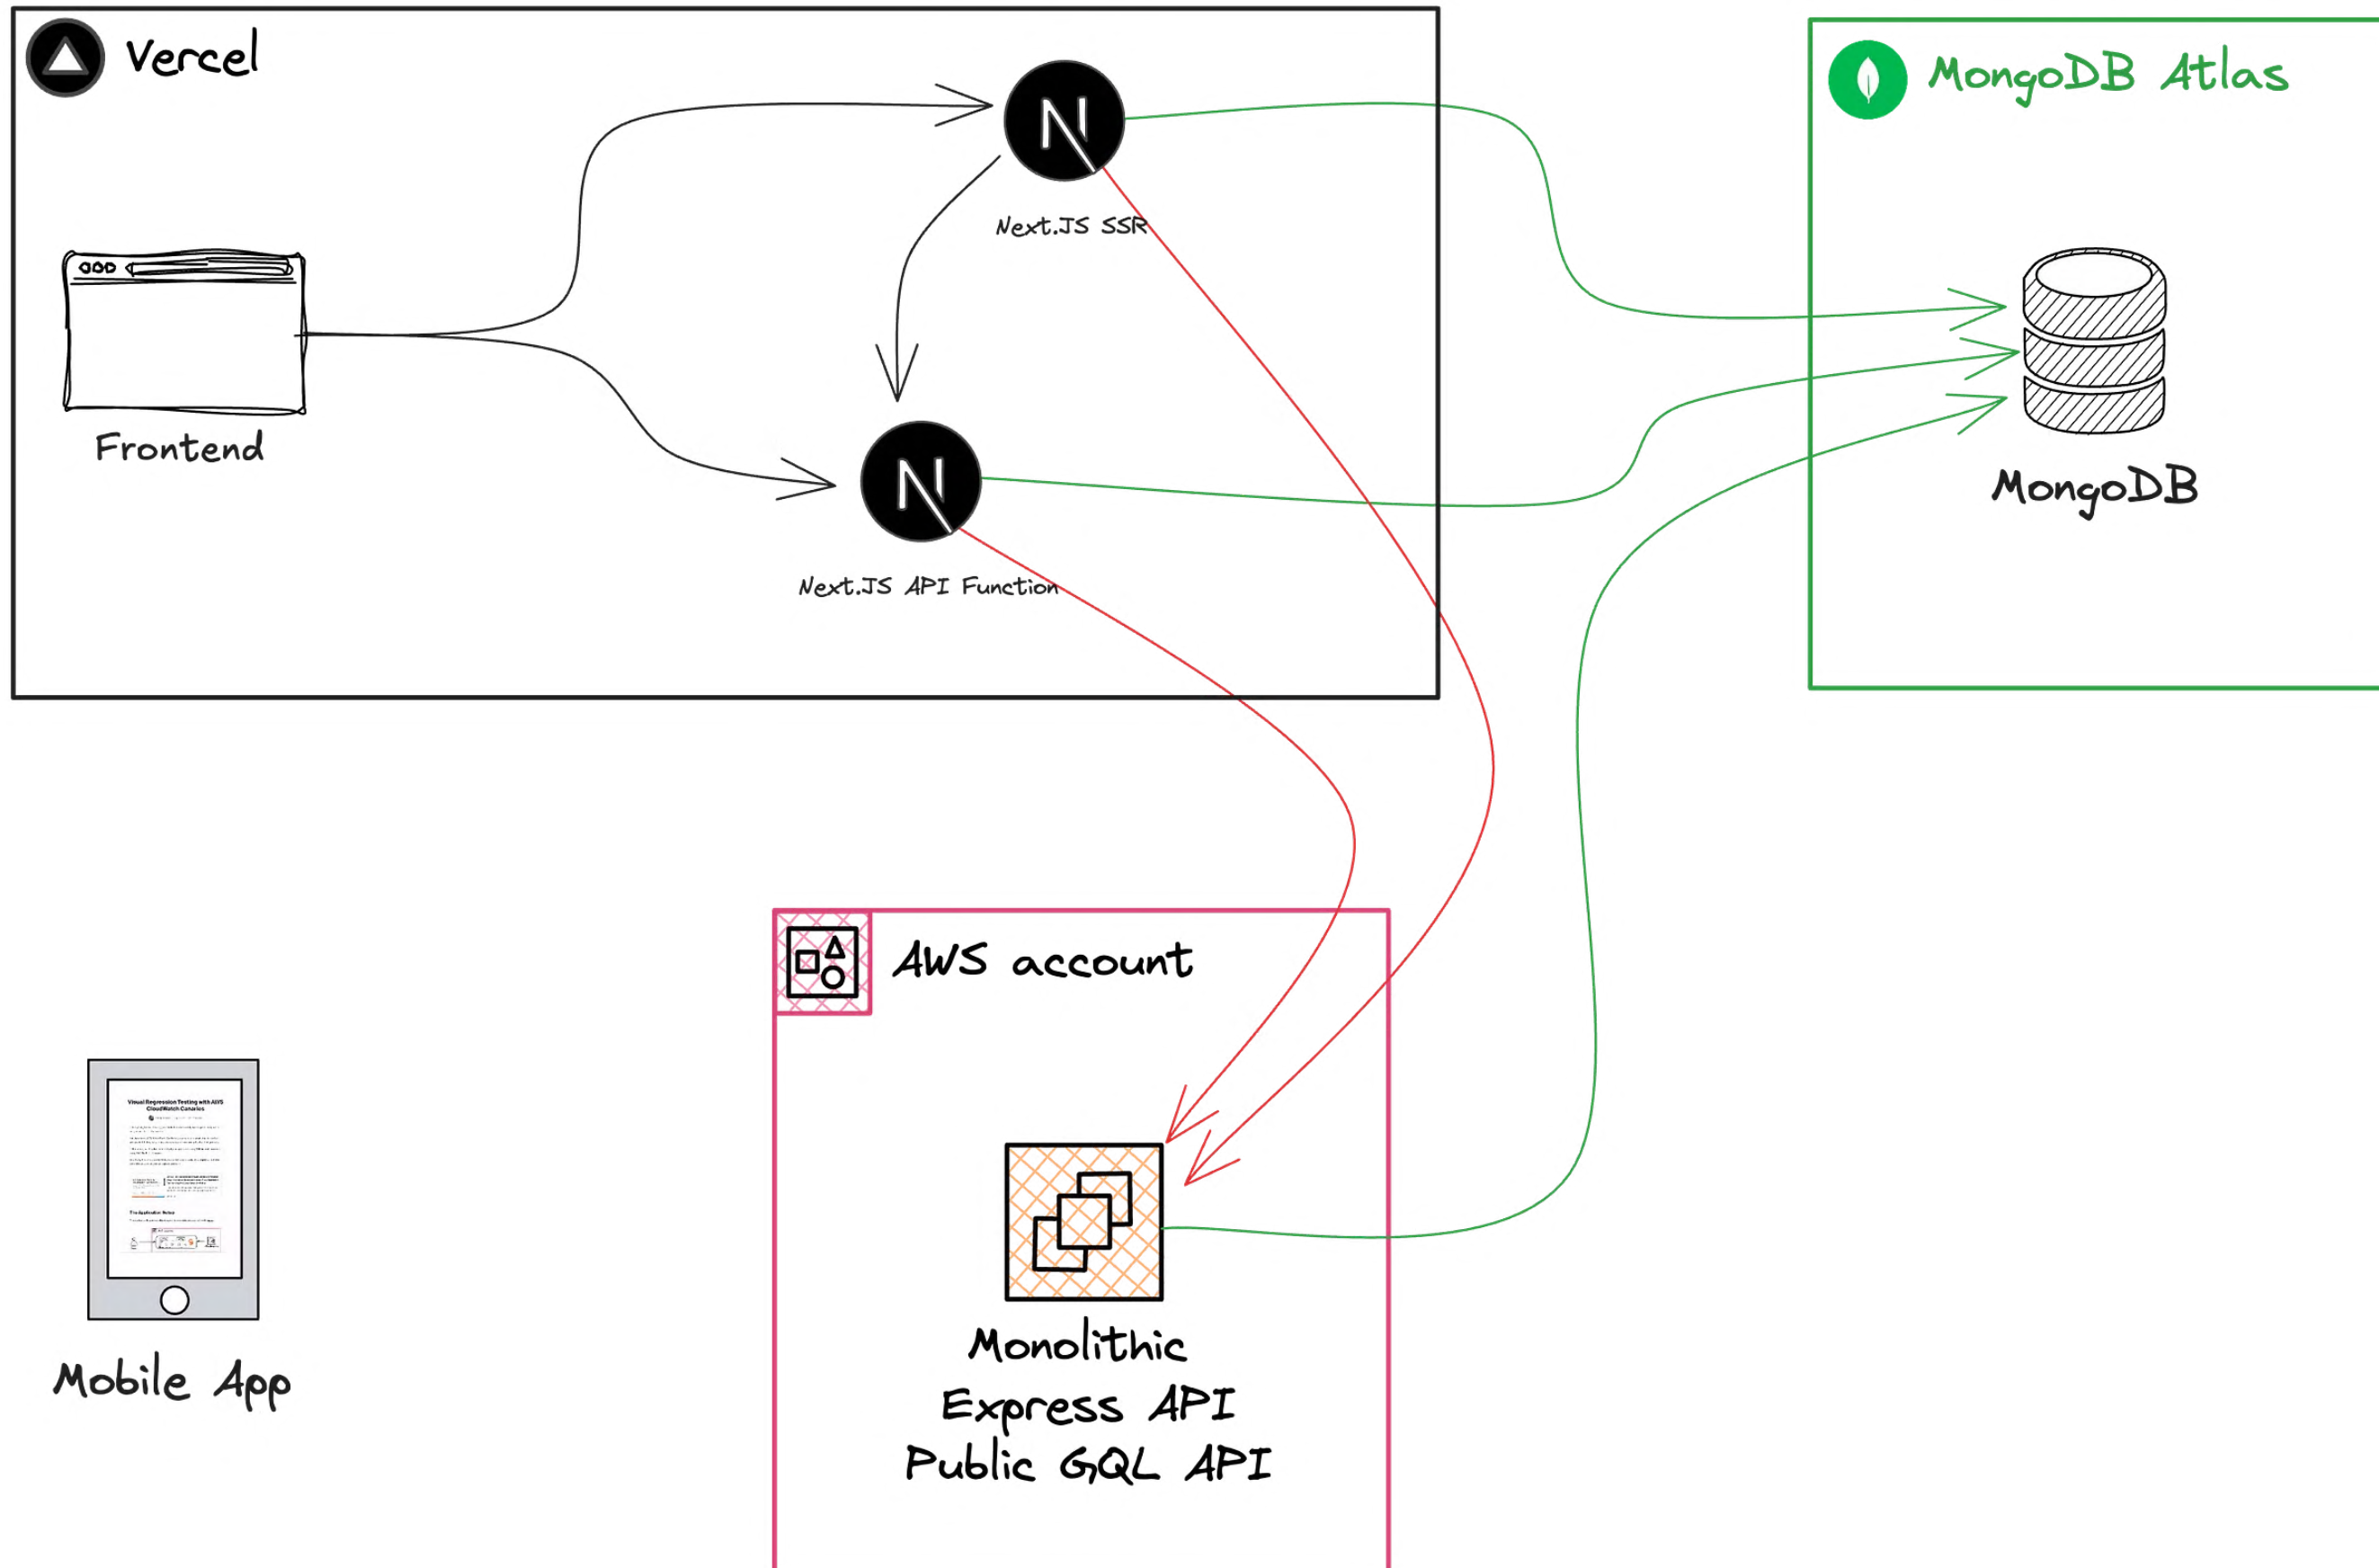


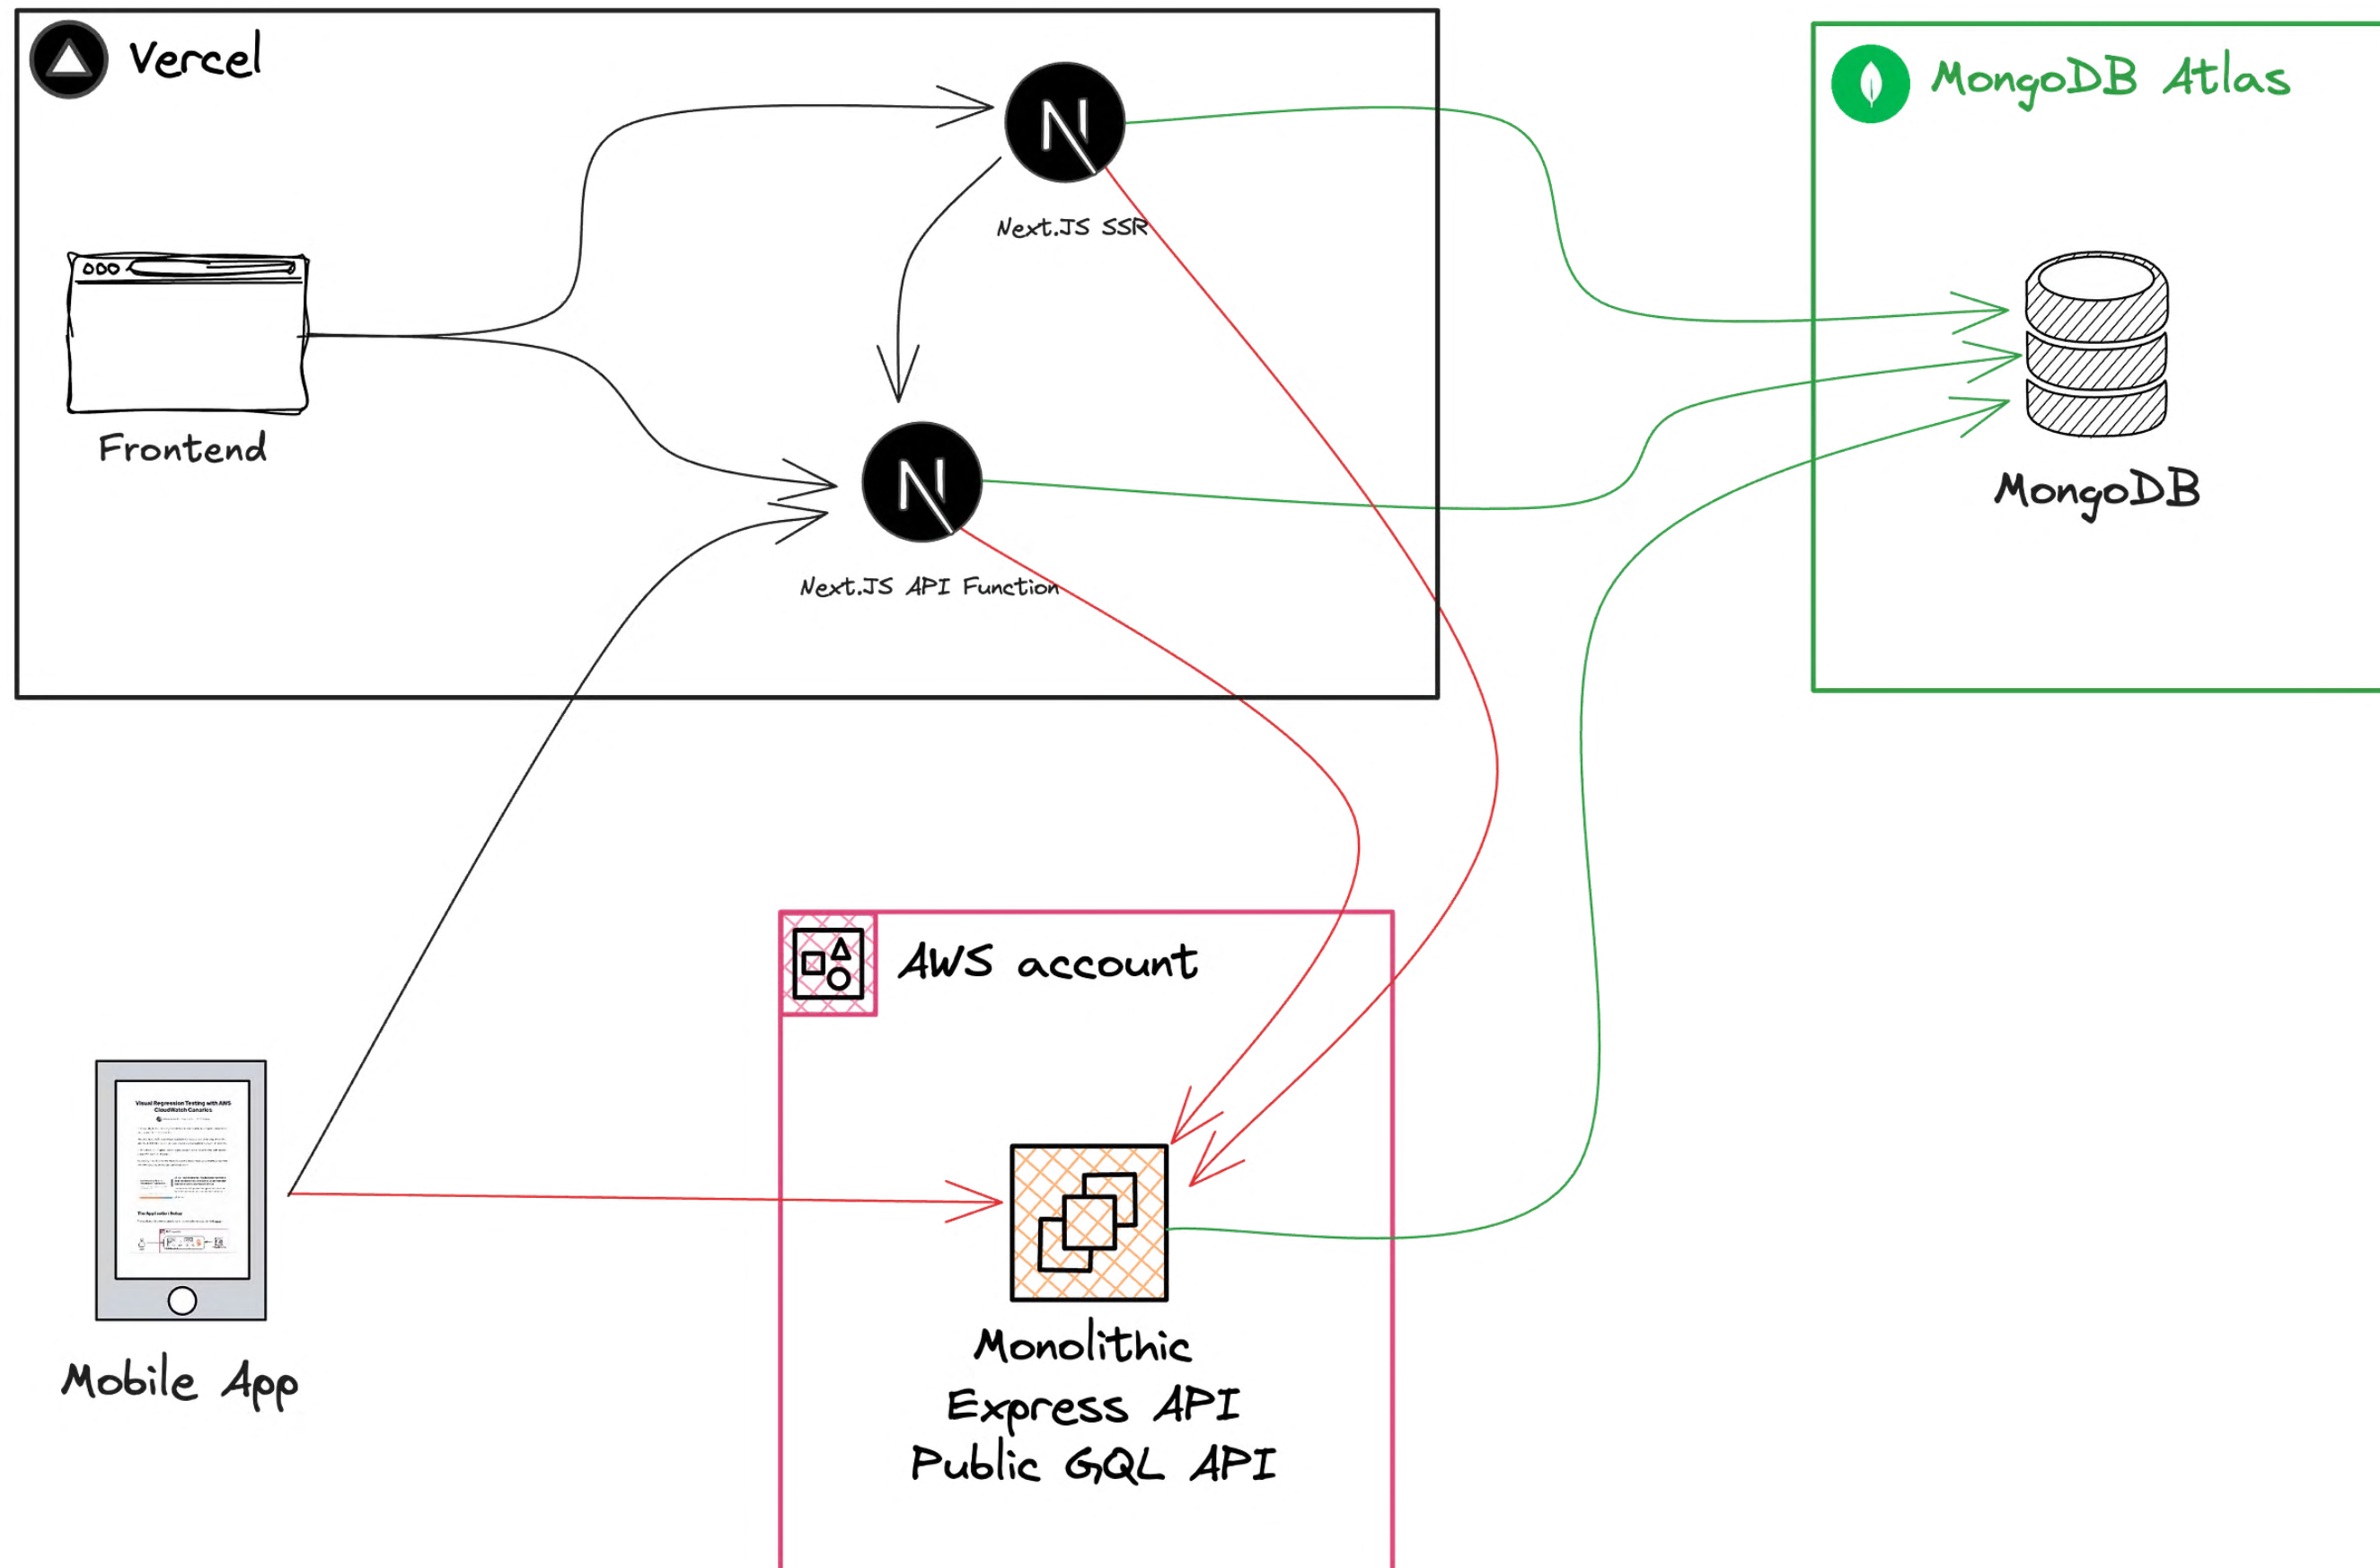












# Downsides 👎

## Standardization

- Business Logic mixed with Data Access
- Authentication always done differently
- Observability (VM, Sentry, CloudWatch, Axiom, Vercel)

## Caching

- Cache only time-based, no purging possible
- Some manual redis caches - no overview


## Developer Experience

- Duplication of API/Implementations
- Type-Safety

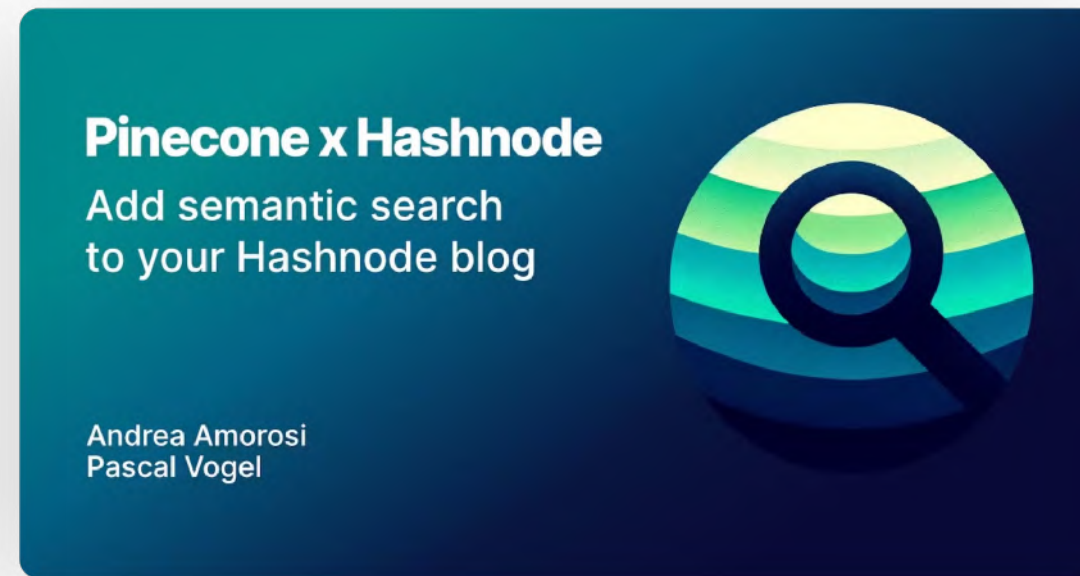
→ **One Central API**

# Why a public API?

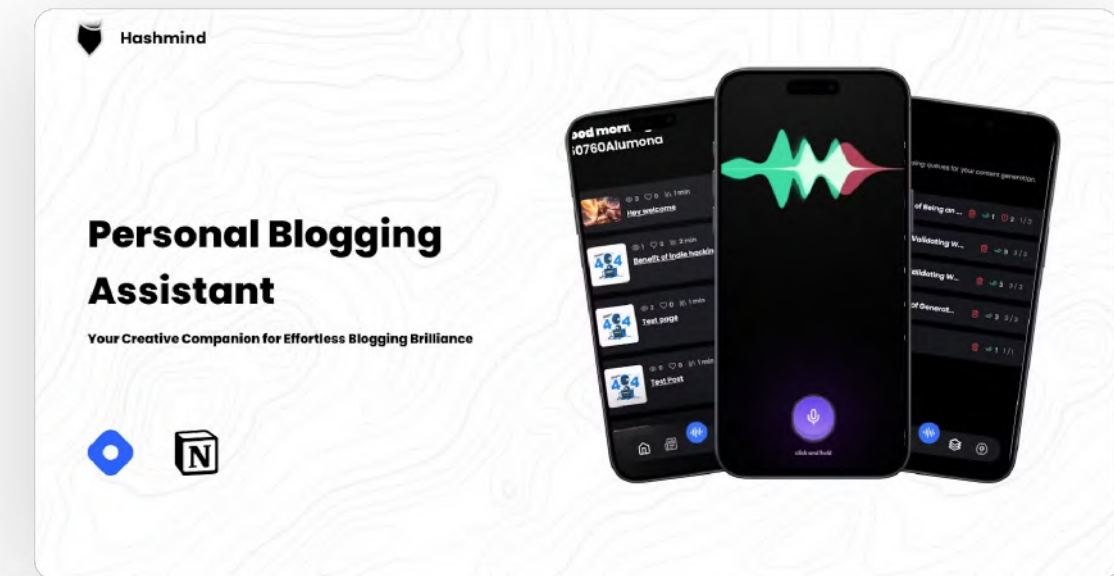
- Our API is in its nature public → Hosting of public blog posts
- Give developers more freedom to build integrations



*[Allen Helton, Serverless Cross-Post](#)*



*[Andrea Armosi, Pascal Vogel - Semantic Search](#)*



*[Benaiah - Hashmind Blogging Assistant](#)*

→ **USE HASHNODE AS A HEADLESS CMS**



# Introduction to GraphQL

OPERATIONS | SCHEMA | ONE ENDPOINT



# Introduction to GraphQL

OPERATIONS



```
query User {
  user(username: "SandroVolpicella") {
    id
    username
    bio {
      text
    }
  }
}
```

### Queries

Choose which data to get

```
mutation PublishPost {
  publishPost(
    input: {
      title: "Amazing Post",
      publicationId: "1234",
      contentMarkdown: "# Heading 1"
    }
  ) {
    post {
      id
      slug
    }
  }
}
```

### Mutations

Mutate data. Create posts, remove users, buy pro plan.

```
subscription OnPostPublished {
  postPublished {
    id
    title
    slug
  }
}
```

### Subscriptions

Listen to changes and get real-time updates.



# Introduction to GraphQL

SCHEMA

# The Schema SDL

```
type Publication implements Node {
  """The about section of the publication."""
  about: Content

  """
  Returns the list of drafts of the authenticated user in the publication.
  """
  drafts(
    """A cursor to the last item in the previous page."""
    after: String
    """The filters to be applied to the draft list."""
    filter: PublicationDraftConnectionFilter
    """The number of drafts to return."""
    first: Int! @constraint(min: 1, max: 50)
  ): DraftConnection! @requireAuth





  isTeam: Boolean!

  pro: PublicationProInfo @hidden

  scheduledDrafts(
    """A cursor to the last item in the previous page."""
    after: String
    """The filters to be applied to the draft list."""
    filter: PublicationDraftConnectionFilter
    """The number of scheduled drafts to return."""
    first: Int! @constraint(min: 1, max: 50)
  ): DraftConnection! @requireAuth
```

# The Schema SDL

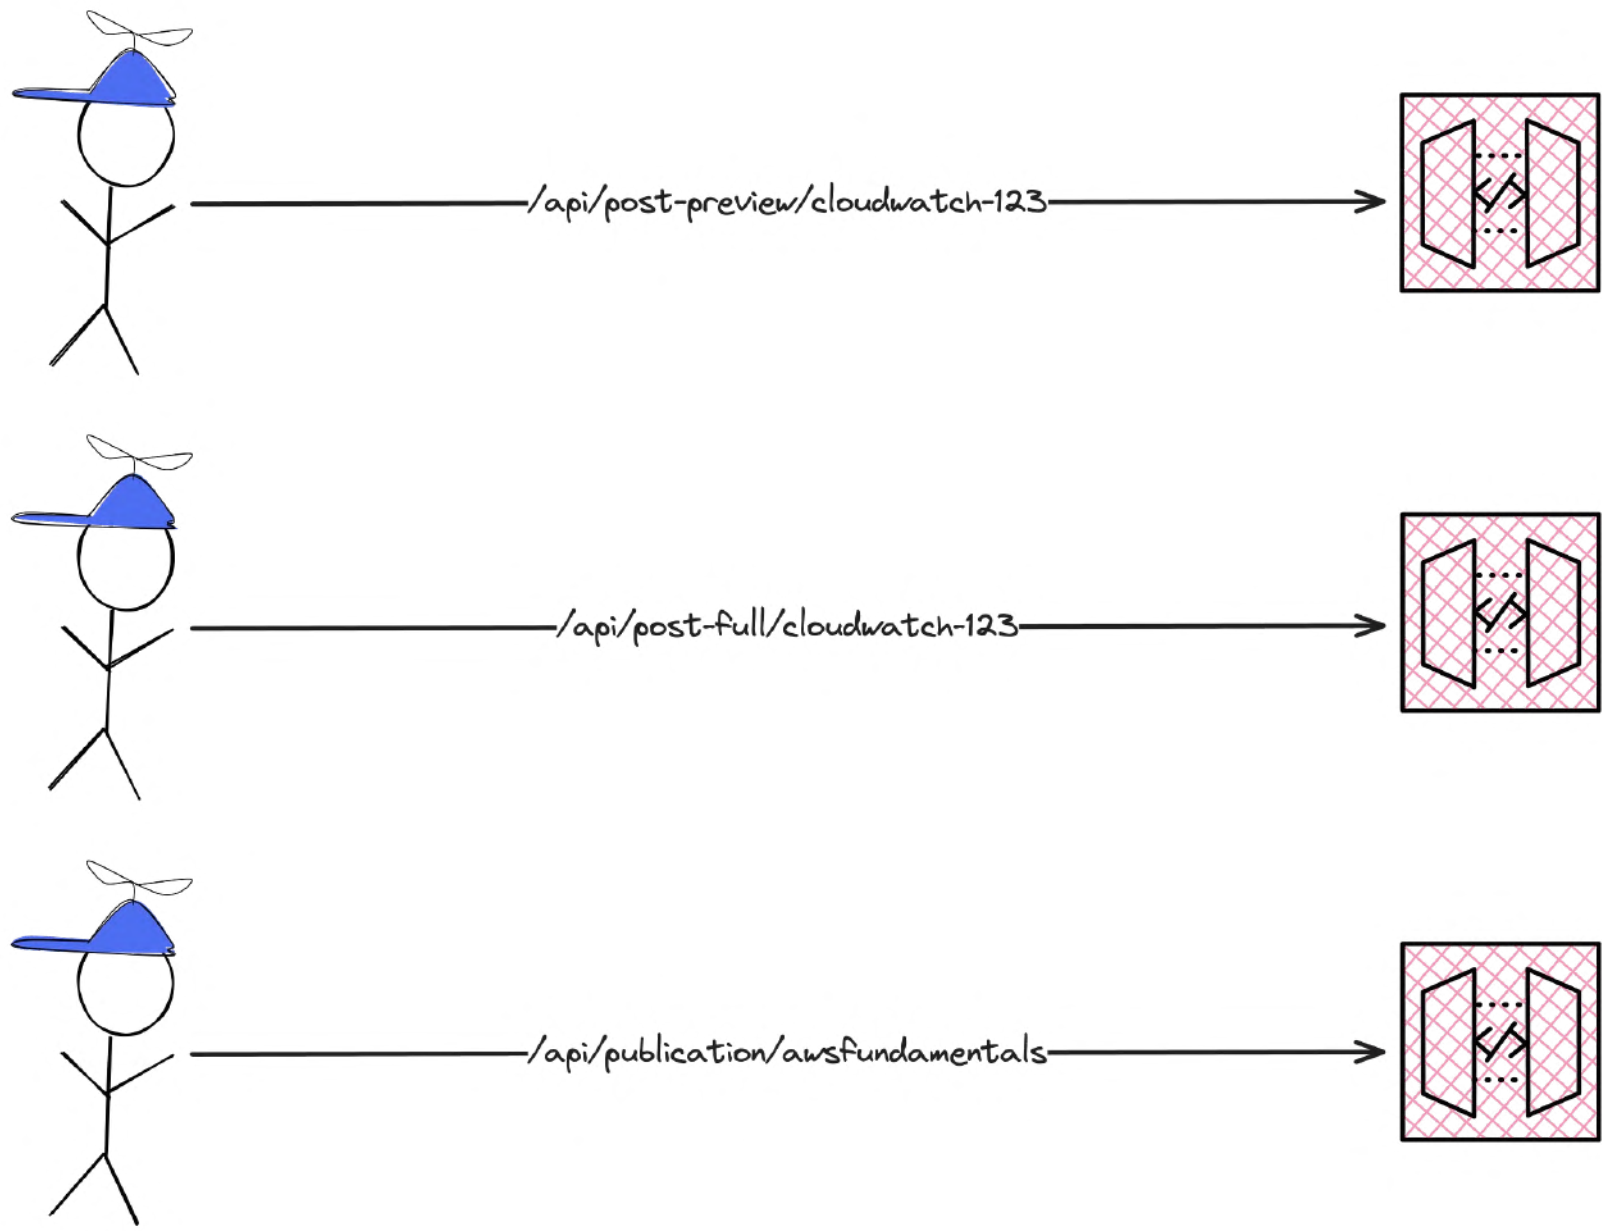
```
type Publication implements Node {  
  """The about section of the publication."""  
  about: Content  
  
  """  
  Returns the list of drafts of the authenticated user in the publication.  
  """  
  drafts(  
    """A cursor to the last item in the previous page."""  
    after: String  
    """The filters to be applied to the draft list."""  
    filter: PublicationDraftConnectionFilter  
    """The number of drafts to return."""  
    first: Int! @constraint(min: 1, max: 50)  
  ): DraftConnection! @requireAuth  
  
  isTeam: Boolean!  
  
  pro: PublicationProInfo @hidden  
  
  scheduledDrafts(  
    """A cursor to the last item in the previous page."""  
    after: String  
    """The filters to be applied to the draft list."""  
    filter: PublicationDraftConnectionFilter  
    """The number of scheduled drafts to return."""  
    first: Int! @constraint(min: 1, max: 50)  
  ): DraftConnection! @requireAuth  
}
```





# Introduction to GraphQL

ONE ENDPOINT



### **REST: Multiple Endpoints**

In REST you need to define several endpoints for similar data



### **GQL: One Endpoint**

GraphQL has one POST Endpoint  
fewer network requests, simpler observability



# Why GraphQL?

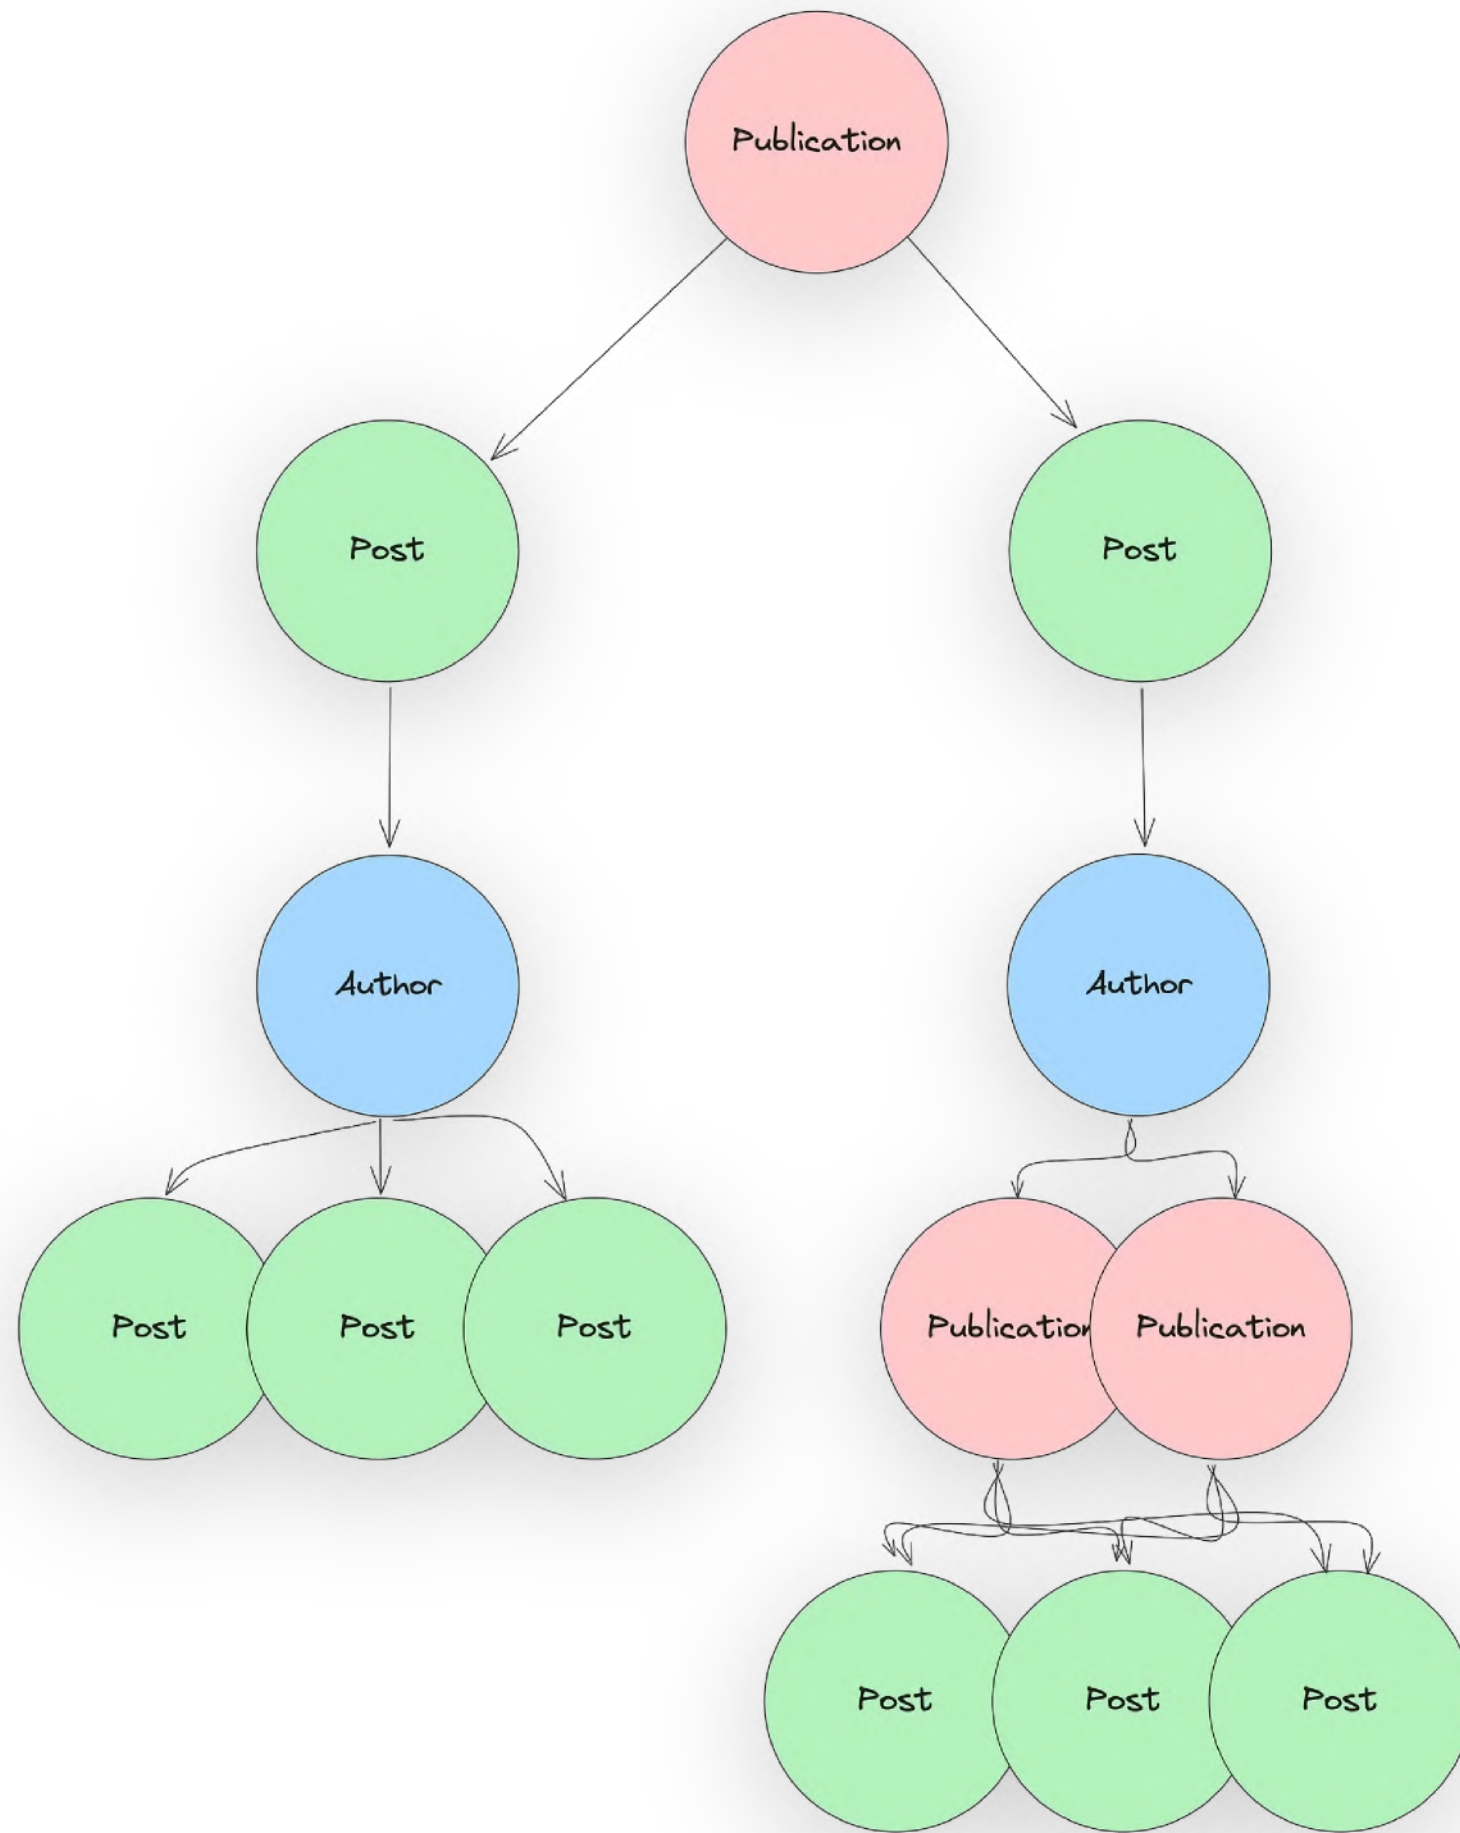
GRAPH | CLIENT CHOOSES DATA | OBSERVABILITY

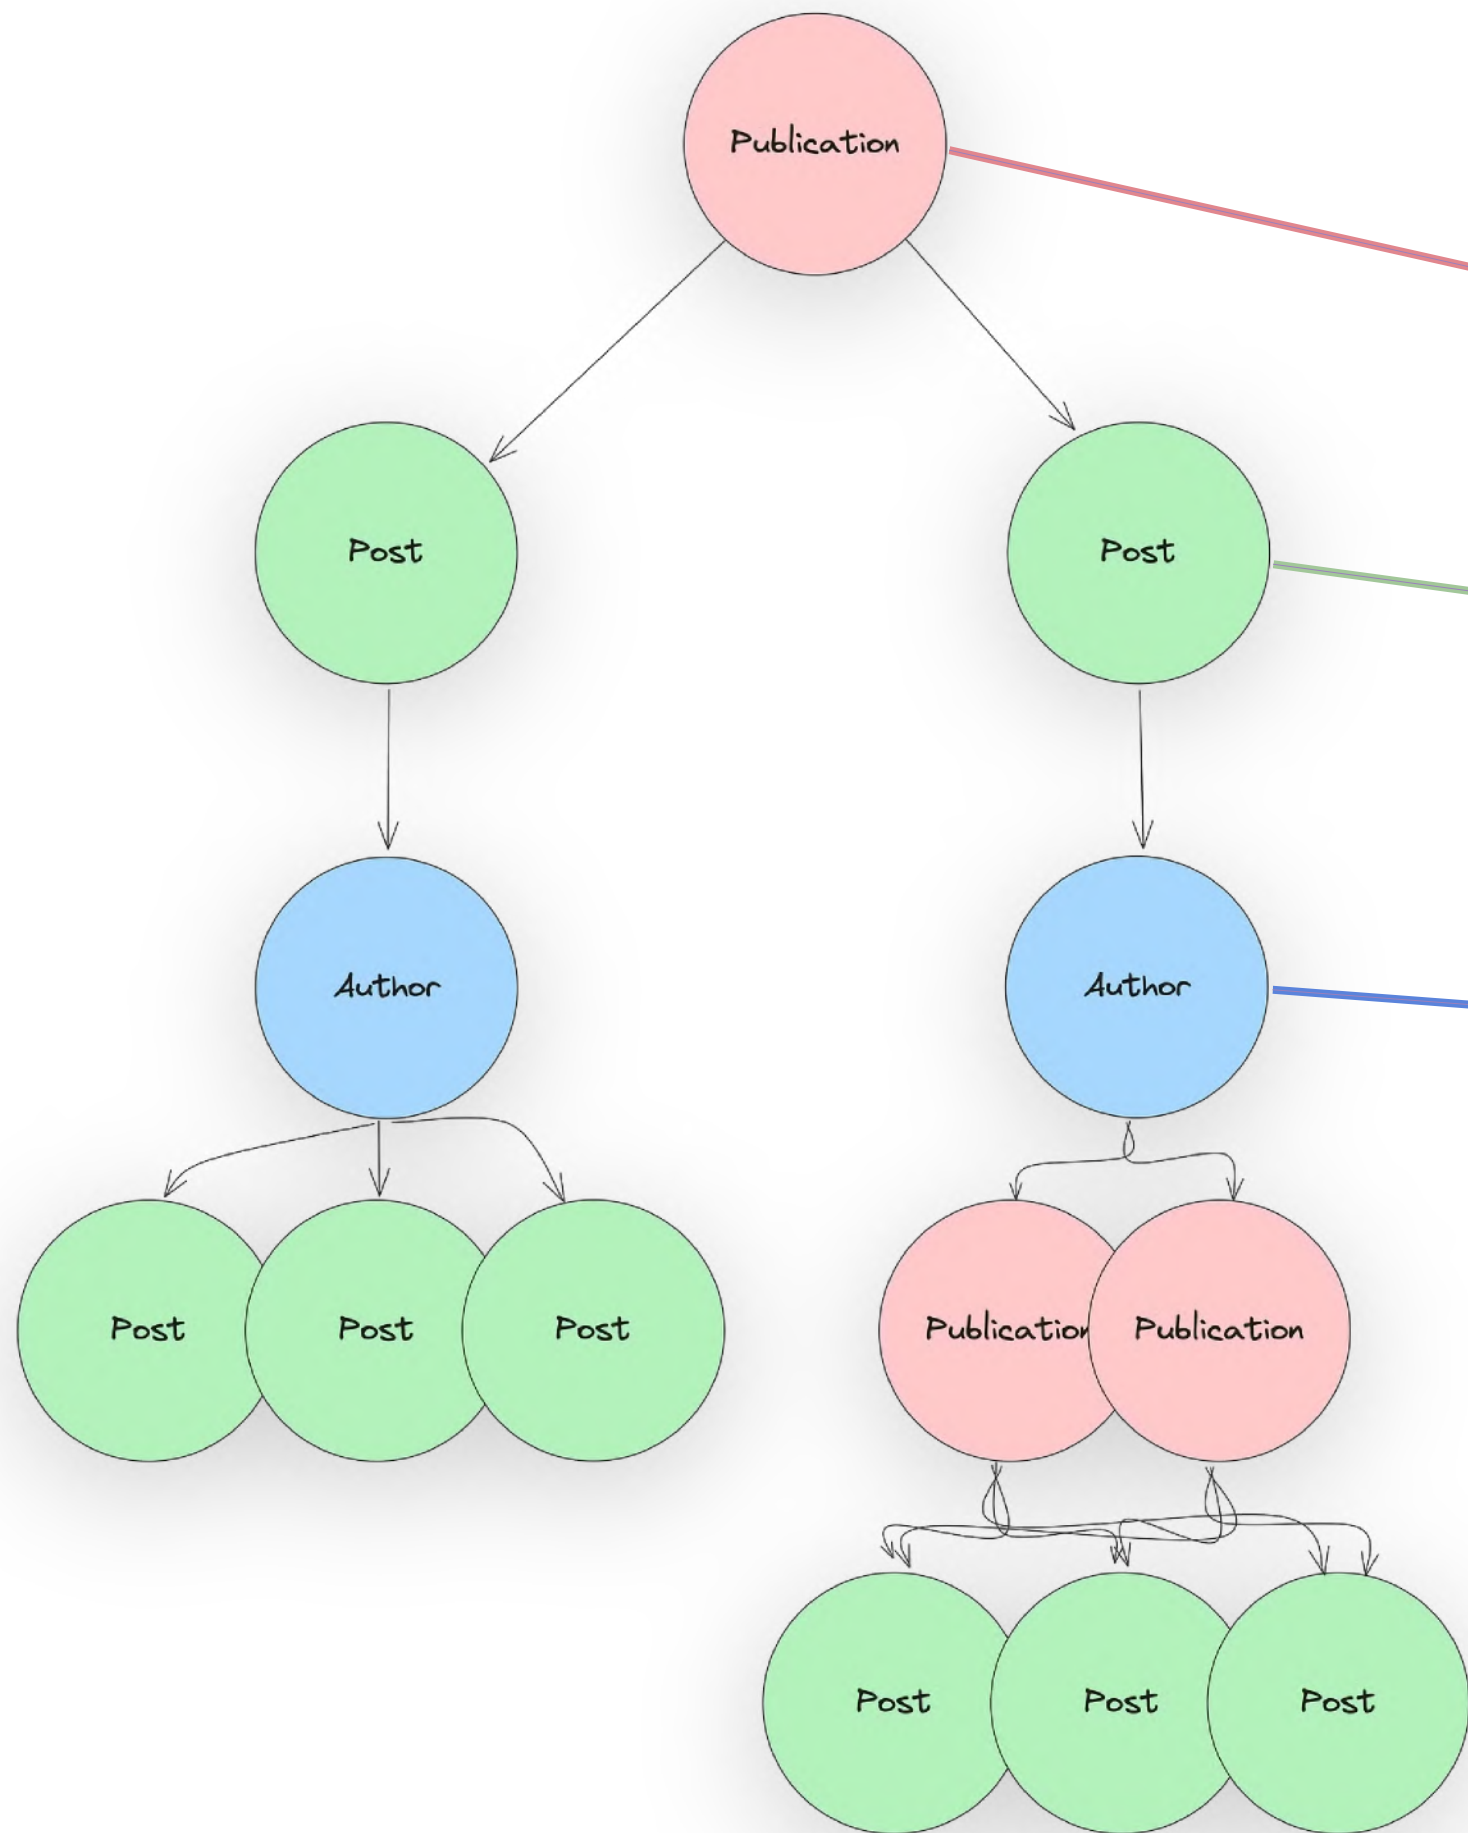


# Why GraphQL?

GRAPH





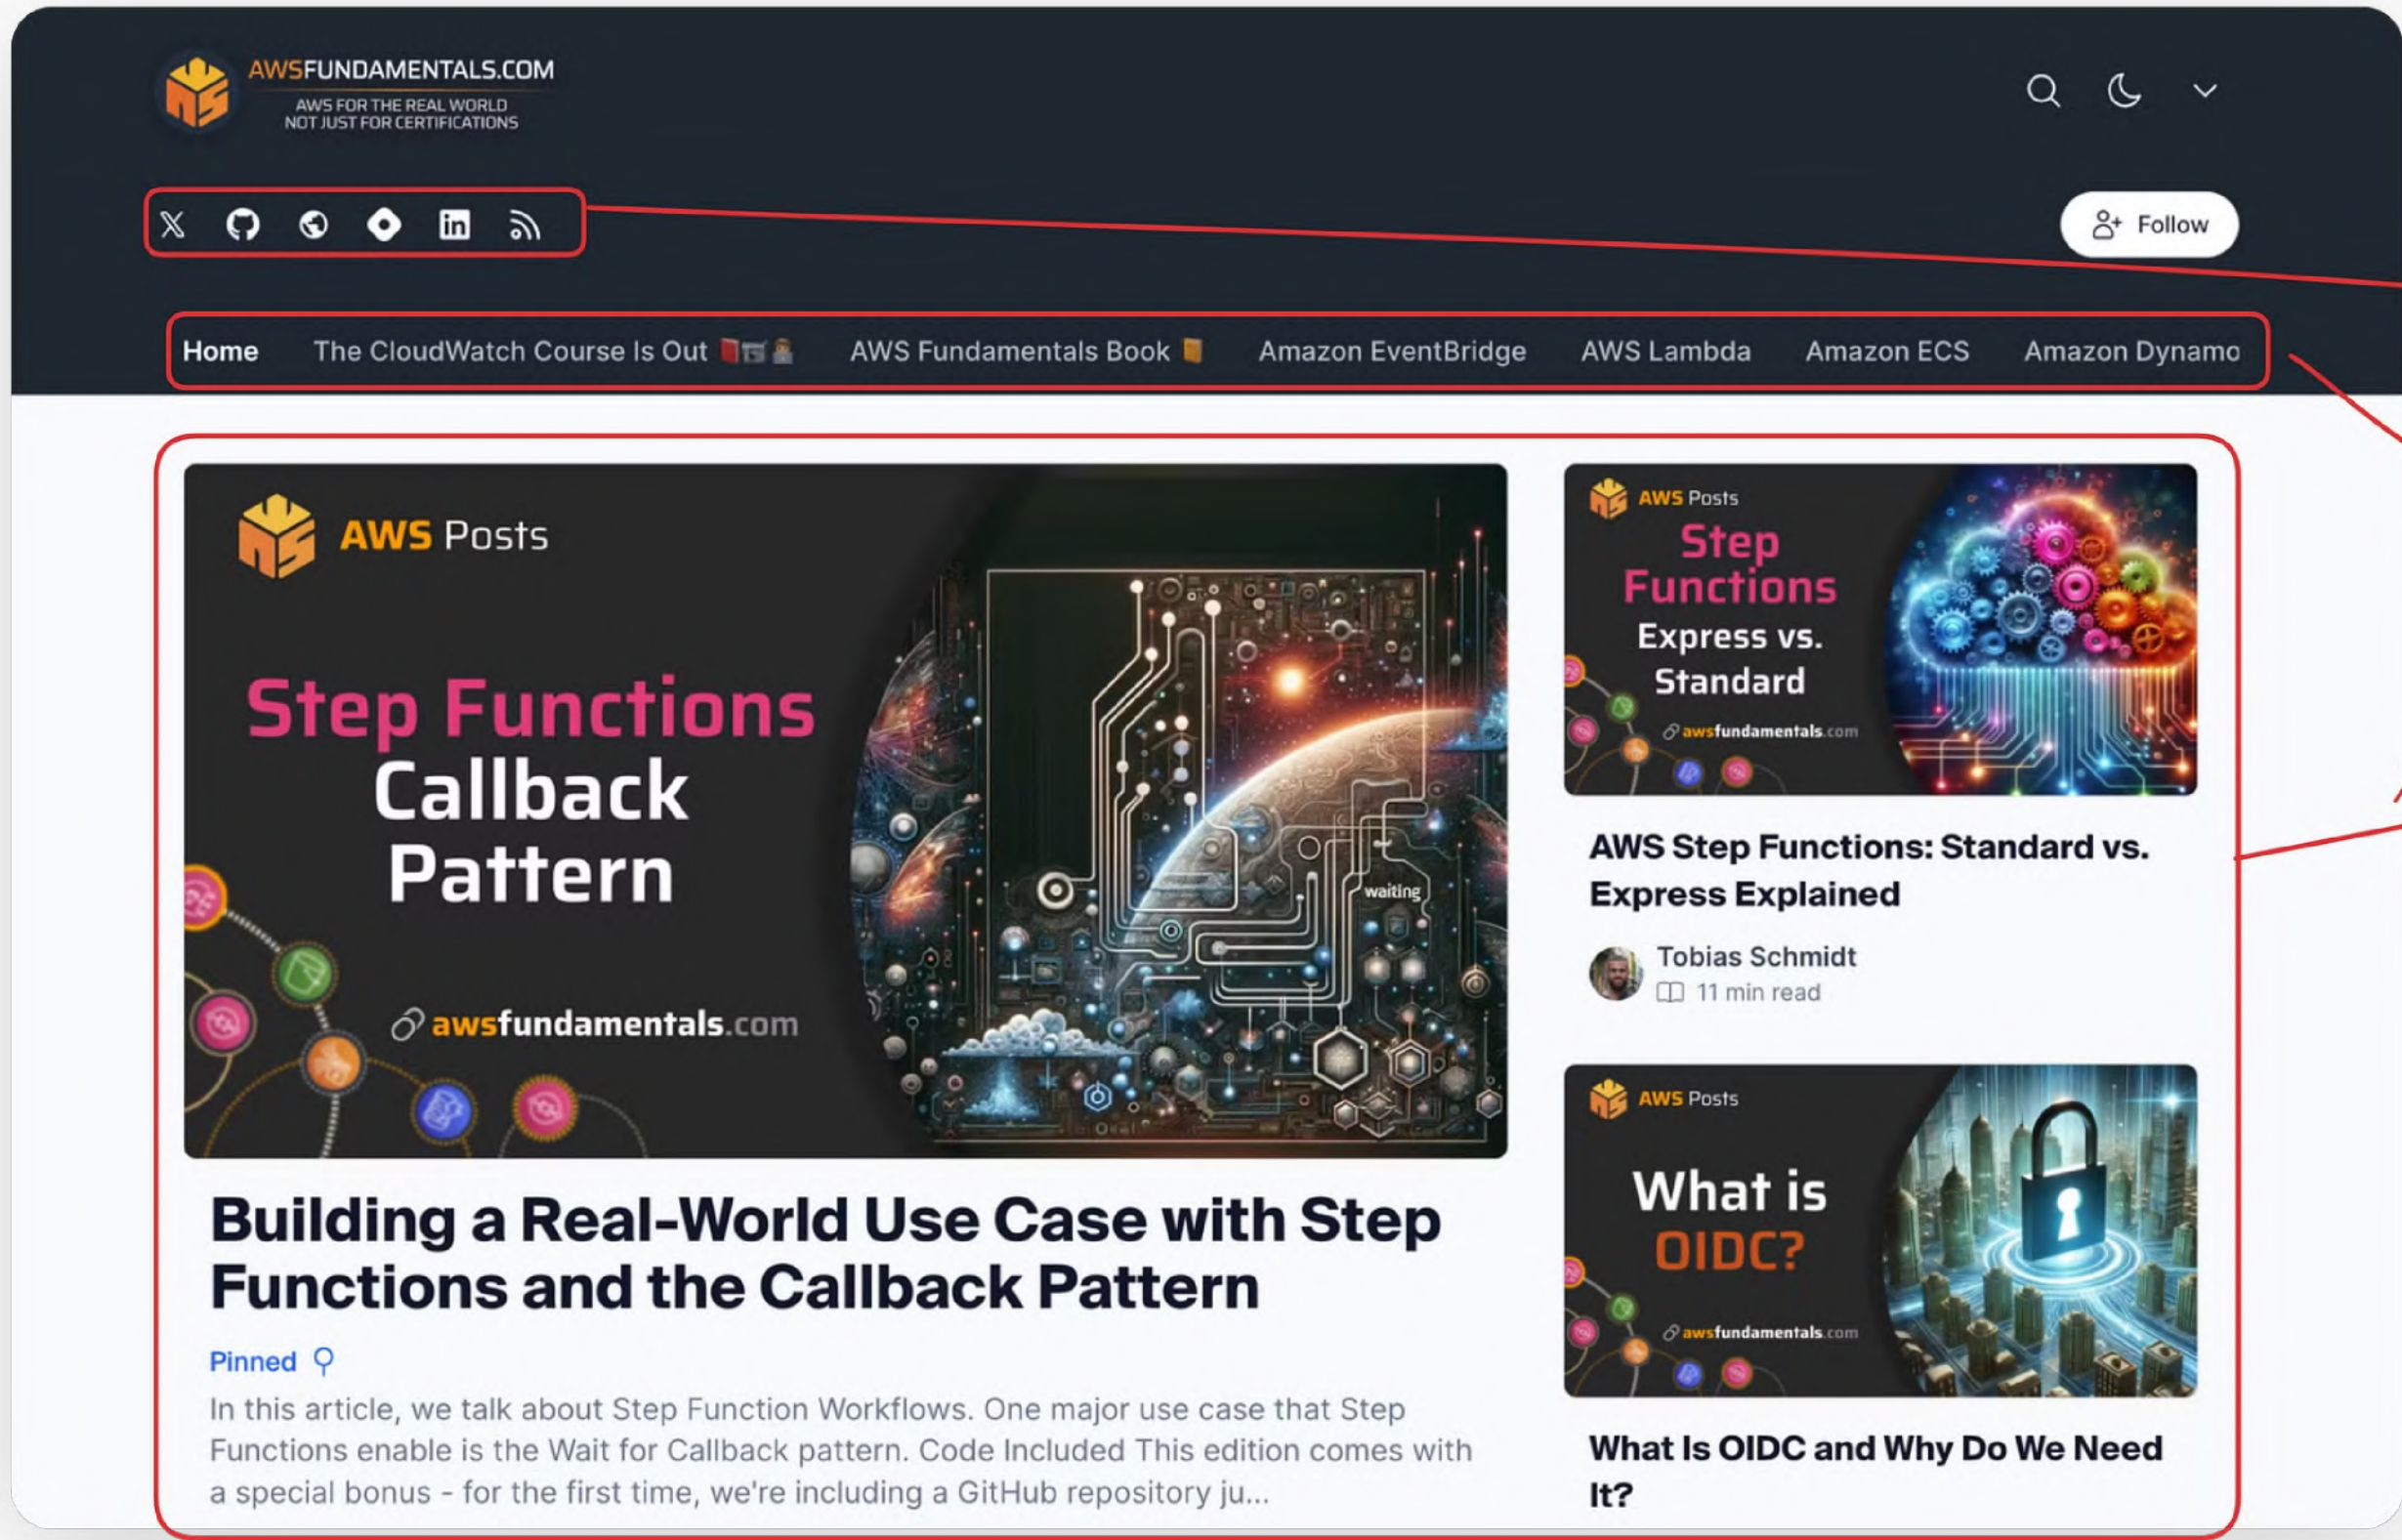


```
query GetPosts {  
  publication(host: "blog.awsfundamentals.com") {  
    id  
    title  
    posts(first: 2) {  
      edges {  
        node {  
          id  
          title  
        }  
      }  
      author {  
        id  
        username  
      }  
    }  
  }  
}
```



# Why GraphQL?

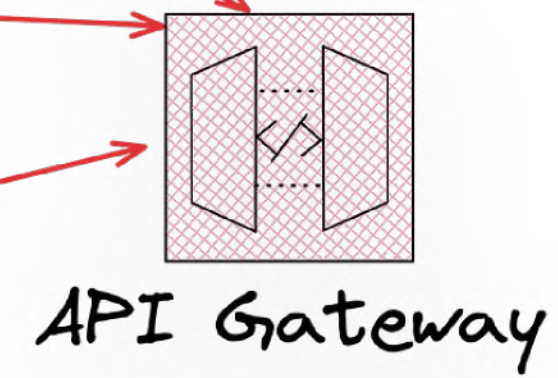
CLIENT CHOOSES DATA

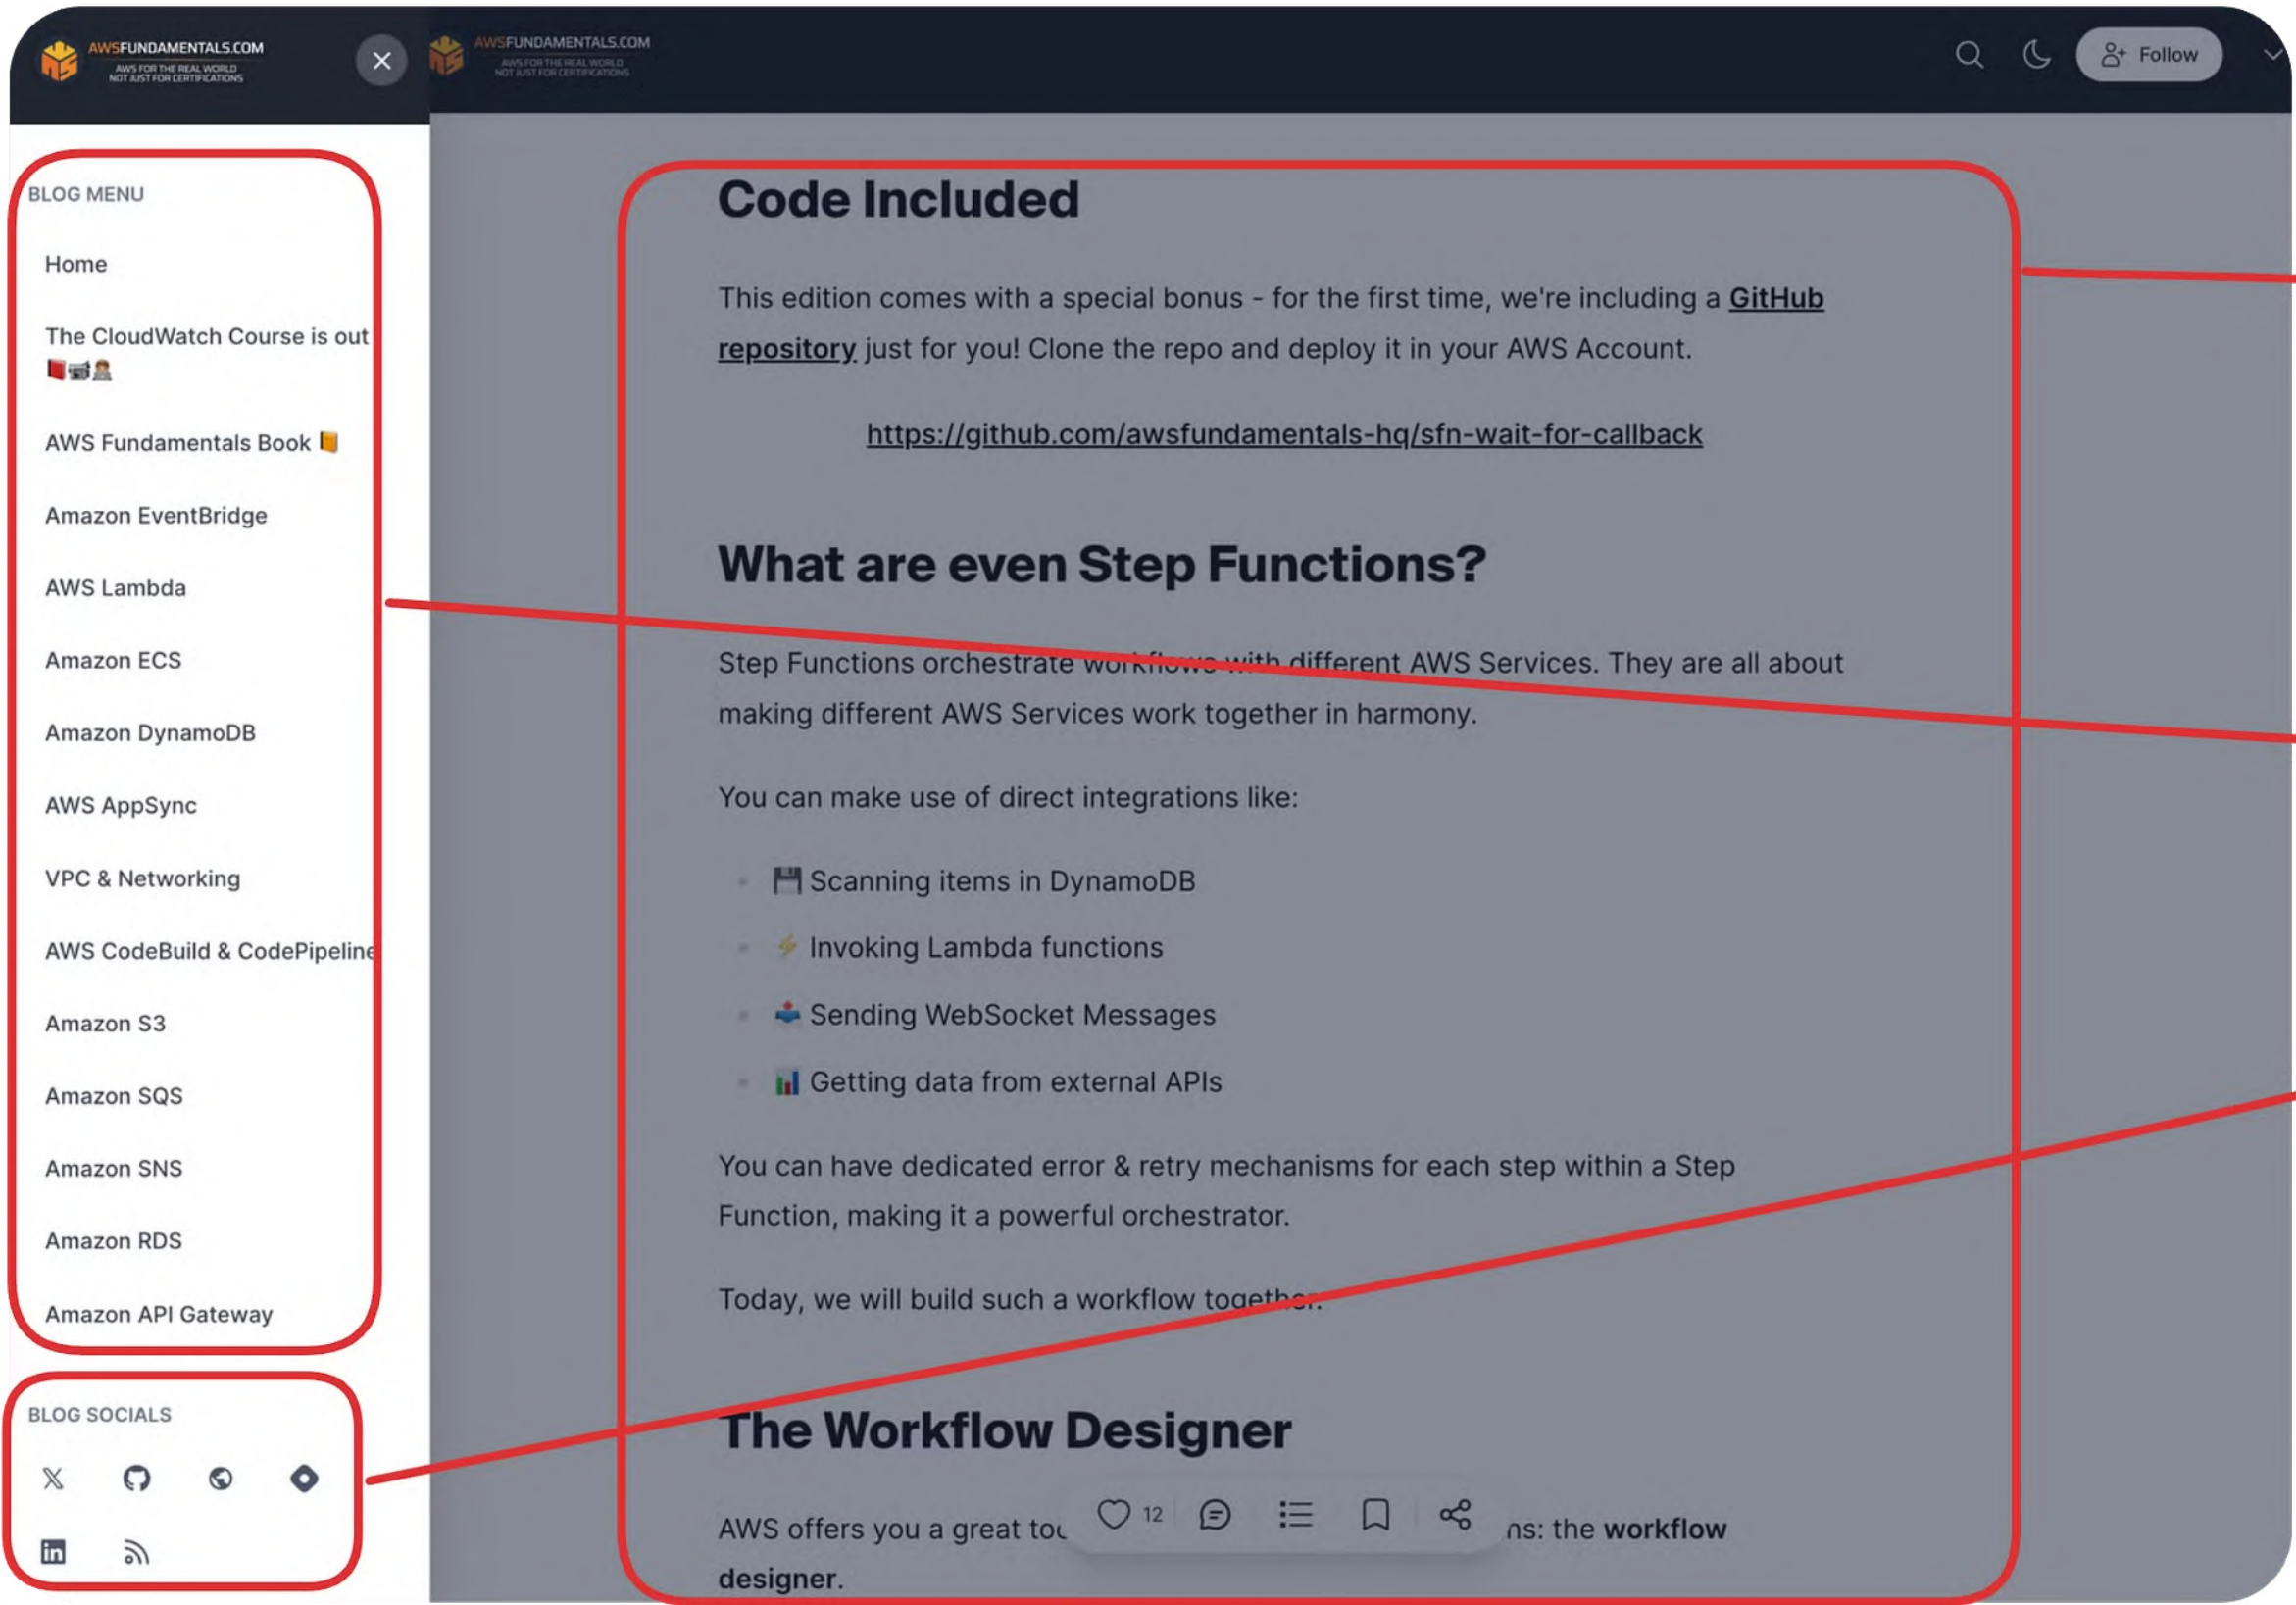


/api/links

/api/navbar

/api/latest-posts






*/api/full-post*

*/api/navbar*

*/api/links*



*API Gateway*

# Why GraphQL over REST → Client chooses data

## PREVIEW

```
query PubWithPreviews {
  publication(host: "blog.awsfundamentals.com") {
    id
    title
    links {
      twitter
    }
    preferences {
      navBarItems {
        id
        label
      }
    }
  }
  posts(first: 3) {
    edges {
      node {
        id
        title
        author {
          name
        }
      }
    }
  }
}
```

## FULL POST

```
query PostWithContent {
  publication(host: "blog.awsfundamentals.com") {
    id
    title
    links {
      twitter
    }
    preferences {
      navBarItems {
        id
        label
      }
    }
  }
  posts(first: 3) {
    edges {
      node {
        id
        title
        author {
          name
        }
        content {
          markdown
        }
      }
    }
  }
}
```



# Why GraphQL?

PLAYGROUND, OBSERVABILITY, TYPE-SAFETY

# Playground, Observability, Type-Safety

**Playground:** Live documentation and execution environment: [gql.hashnode.com](https://gql.hashnode.com)

**Observability:** Errors by types/users

Operations	Types	Fields
1,327	177	955

Most used >	Requests	Highest latency >	P95
SlugCatchAll	38.2M	Publication	20,984ms
HomePageInitial	5.65M	me	13,457ms
TagInitial	4.61M	inv	5,486ms
HomePagePosts	2.4M	Stats	4,606ms
SlugCatchAll	1.94M	MetricsPerBrowser	4,568ms

**Type-Safety:** Code Generation depending on the used queries

The screenshot shows the GraphQL Playground interface. On the left, there is a documentation panel for the `user` query. It states: "Returns the user with the username." and lists the argument `username: String!` with the description "The username of the user to retrieve." On the right, the query editor contains the following query:

```
1 query User{
2   user(username: "SandroVolpicella") {
3     id
4     username
5     publications(first:10) {
6       edges{
7         node{
8           id
9           title
10        }
11      }
12    }
13  }
14 }
```

At the bottom of the editor, there are tabs for "Variables" and "Headers".





# Serverless implementation of GraphQL API

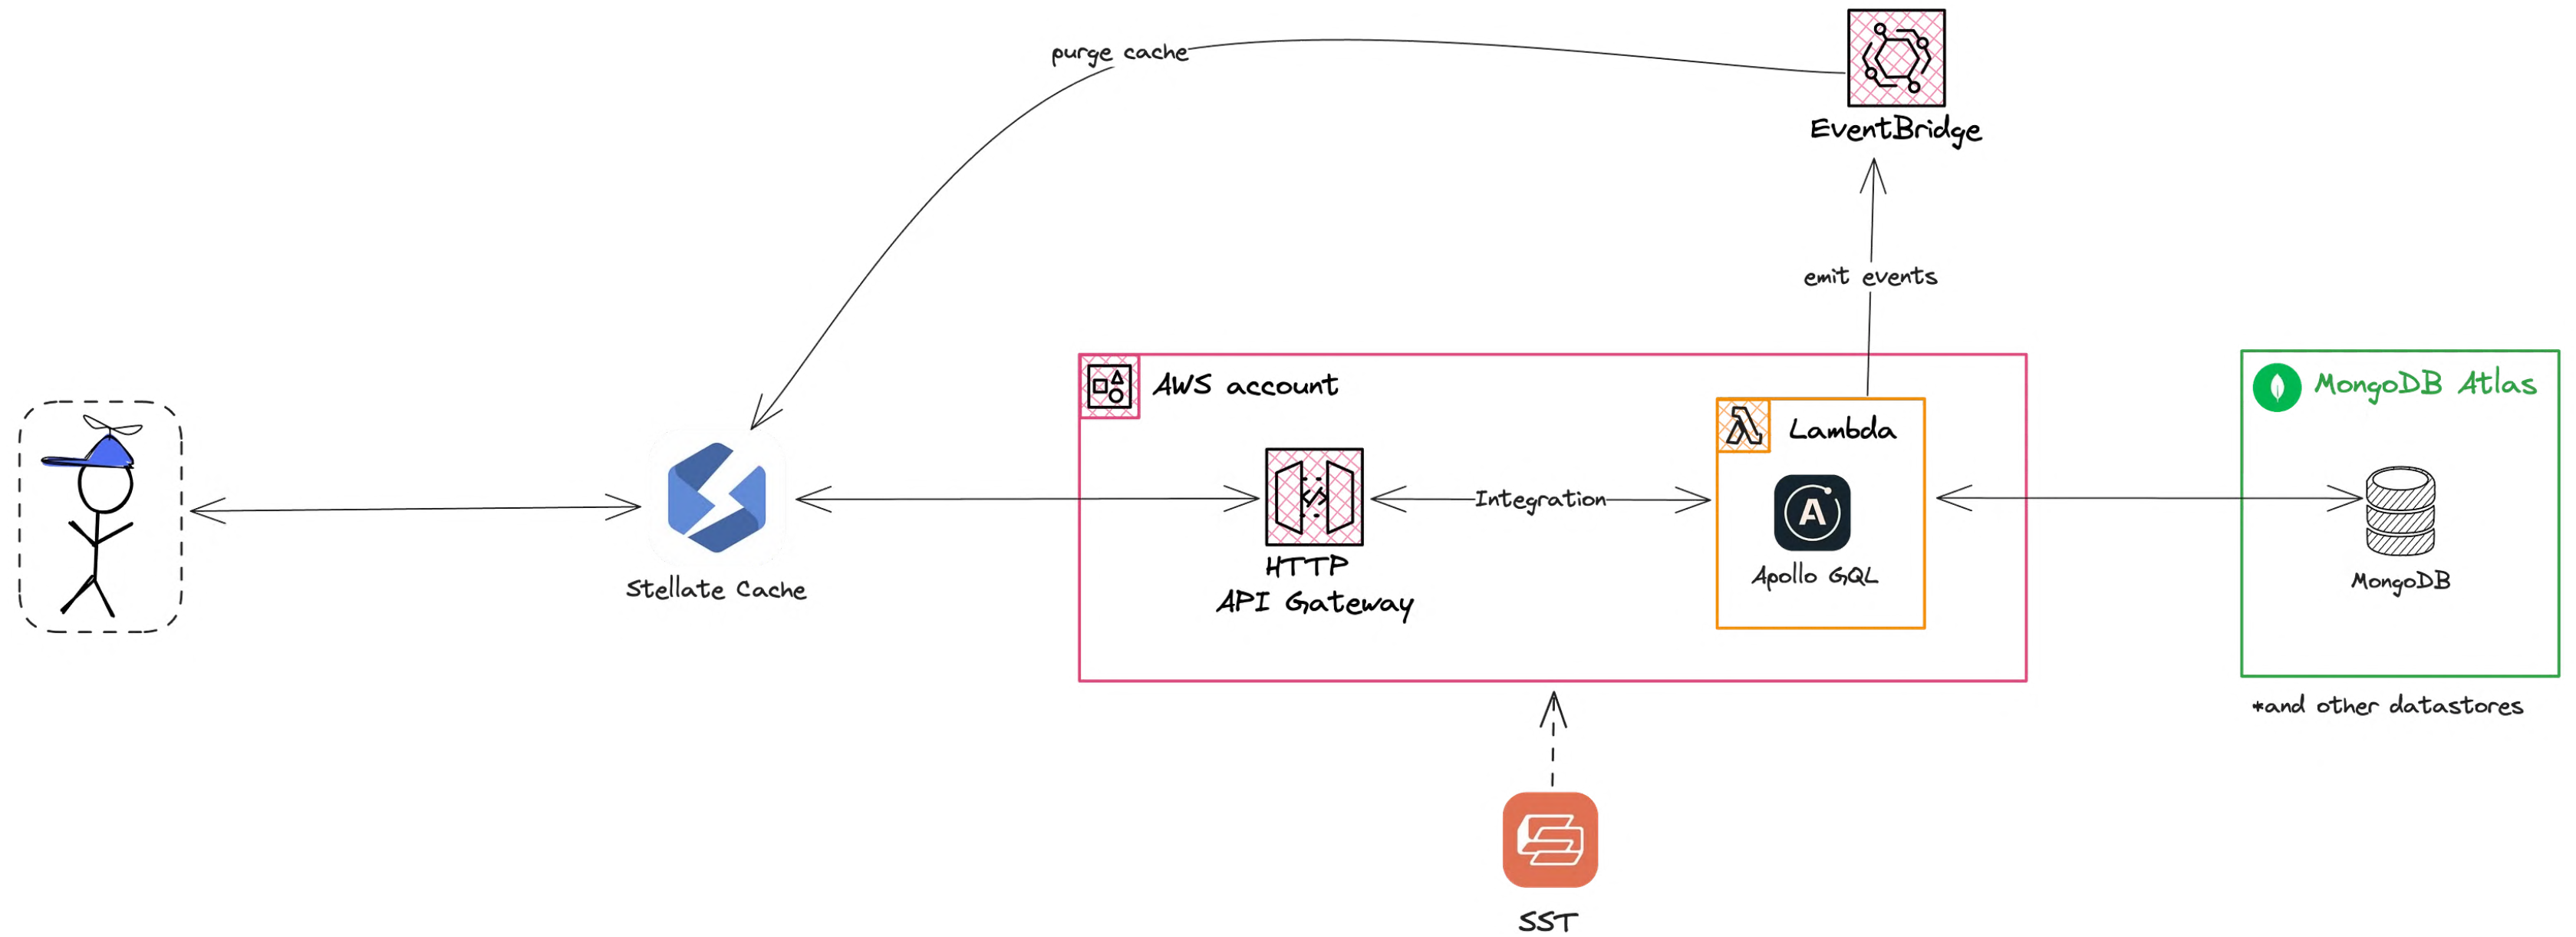
API GATEWAY | LAMBDA | APOLLO | SST



# Serverless implementation of GraphQL API

ARCHITECTURE

# Architecture



# Why no AppSync?

## Technical Misses in AppSync (back then):

- No executable schema
- API should be public → API Key needed (workaround needed)
- Authentication: No cookies in request without workarounds (Lambda at Edge) ⚠️ **now changed**
- VTL hard to build and debug ⚠️ **now JS resolvers**
- We don't need direct integrations

## Main Part:

- Team was familiar with Apollo (old API on Apollo)
- Apollo does have nice features + big community

# Let's talk about cold starts


## Bundle Size

Code properties [Info](#)

Package size  
4 MB

Number of Cold Starts in % (already \*100 🖱️)

0.05



# Let's talk about cold starts

Daily	totalRequests	totalColdStarts	coldStartPct	AverageColdStartMs	P95Ms
2024-06-21...	1455326	1641	0.1128	144.4553	479.9773
2024-06-20...	1296888	931	0.07179	201.8736	657.6846
2024-06-19...	1130920	1008	0.08913	216.9162	722.8642
2024-06-18...	1170211	1043	0.08913	176.4889	617.5303
2024-06-17...	1079429	1157	0.1072	182.4302	657.6846
2024-06-16...	1729176	980	0.05667	108.3641	328.9003
2024-06-15...	1068530	837	0.07833	173.2719	527.5453

- Cold starts **are bad**
- But they **don't occur very often**
- We use **caching a lot**

# Let's t

Daily

2024-06-21...

2024-06-20...

2024-06-19...

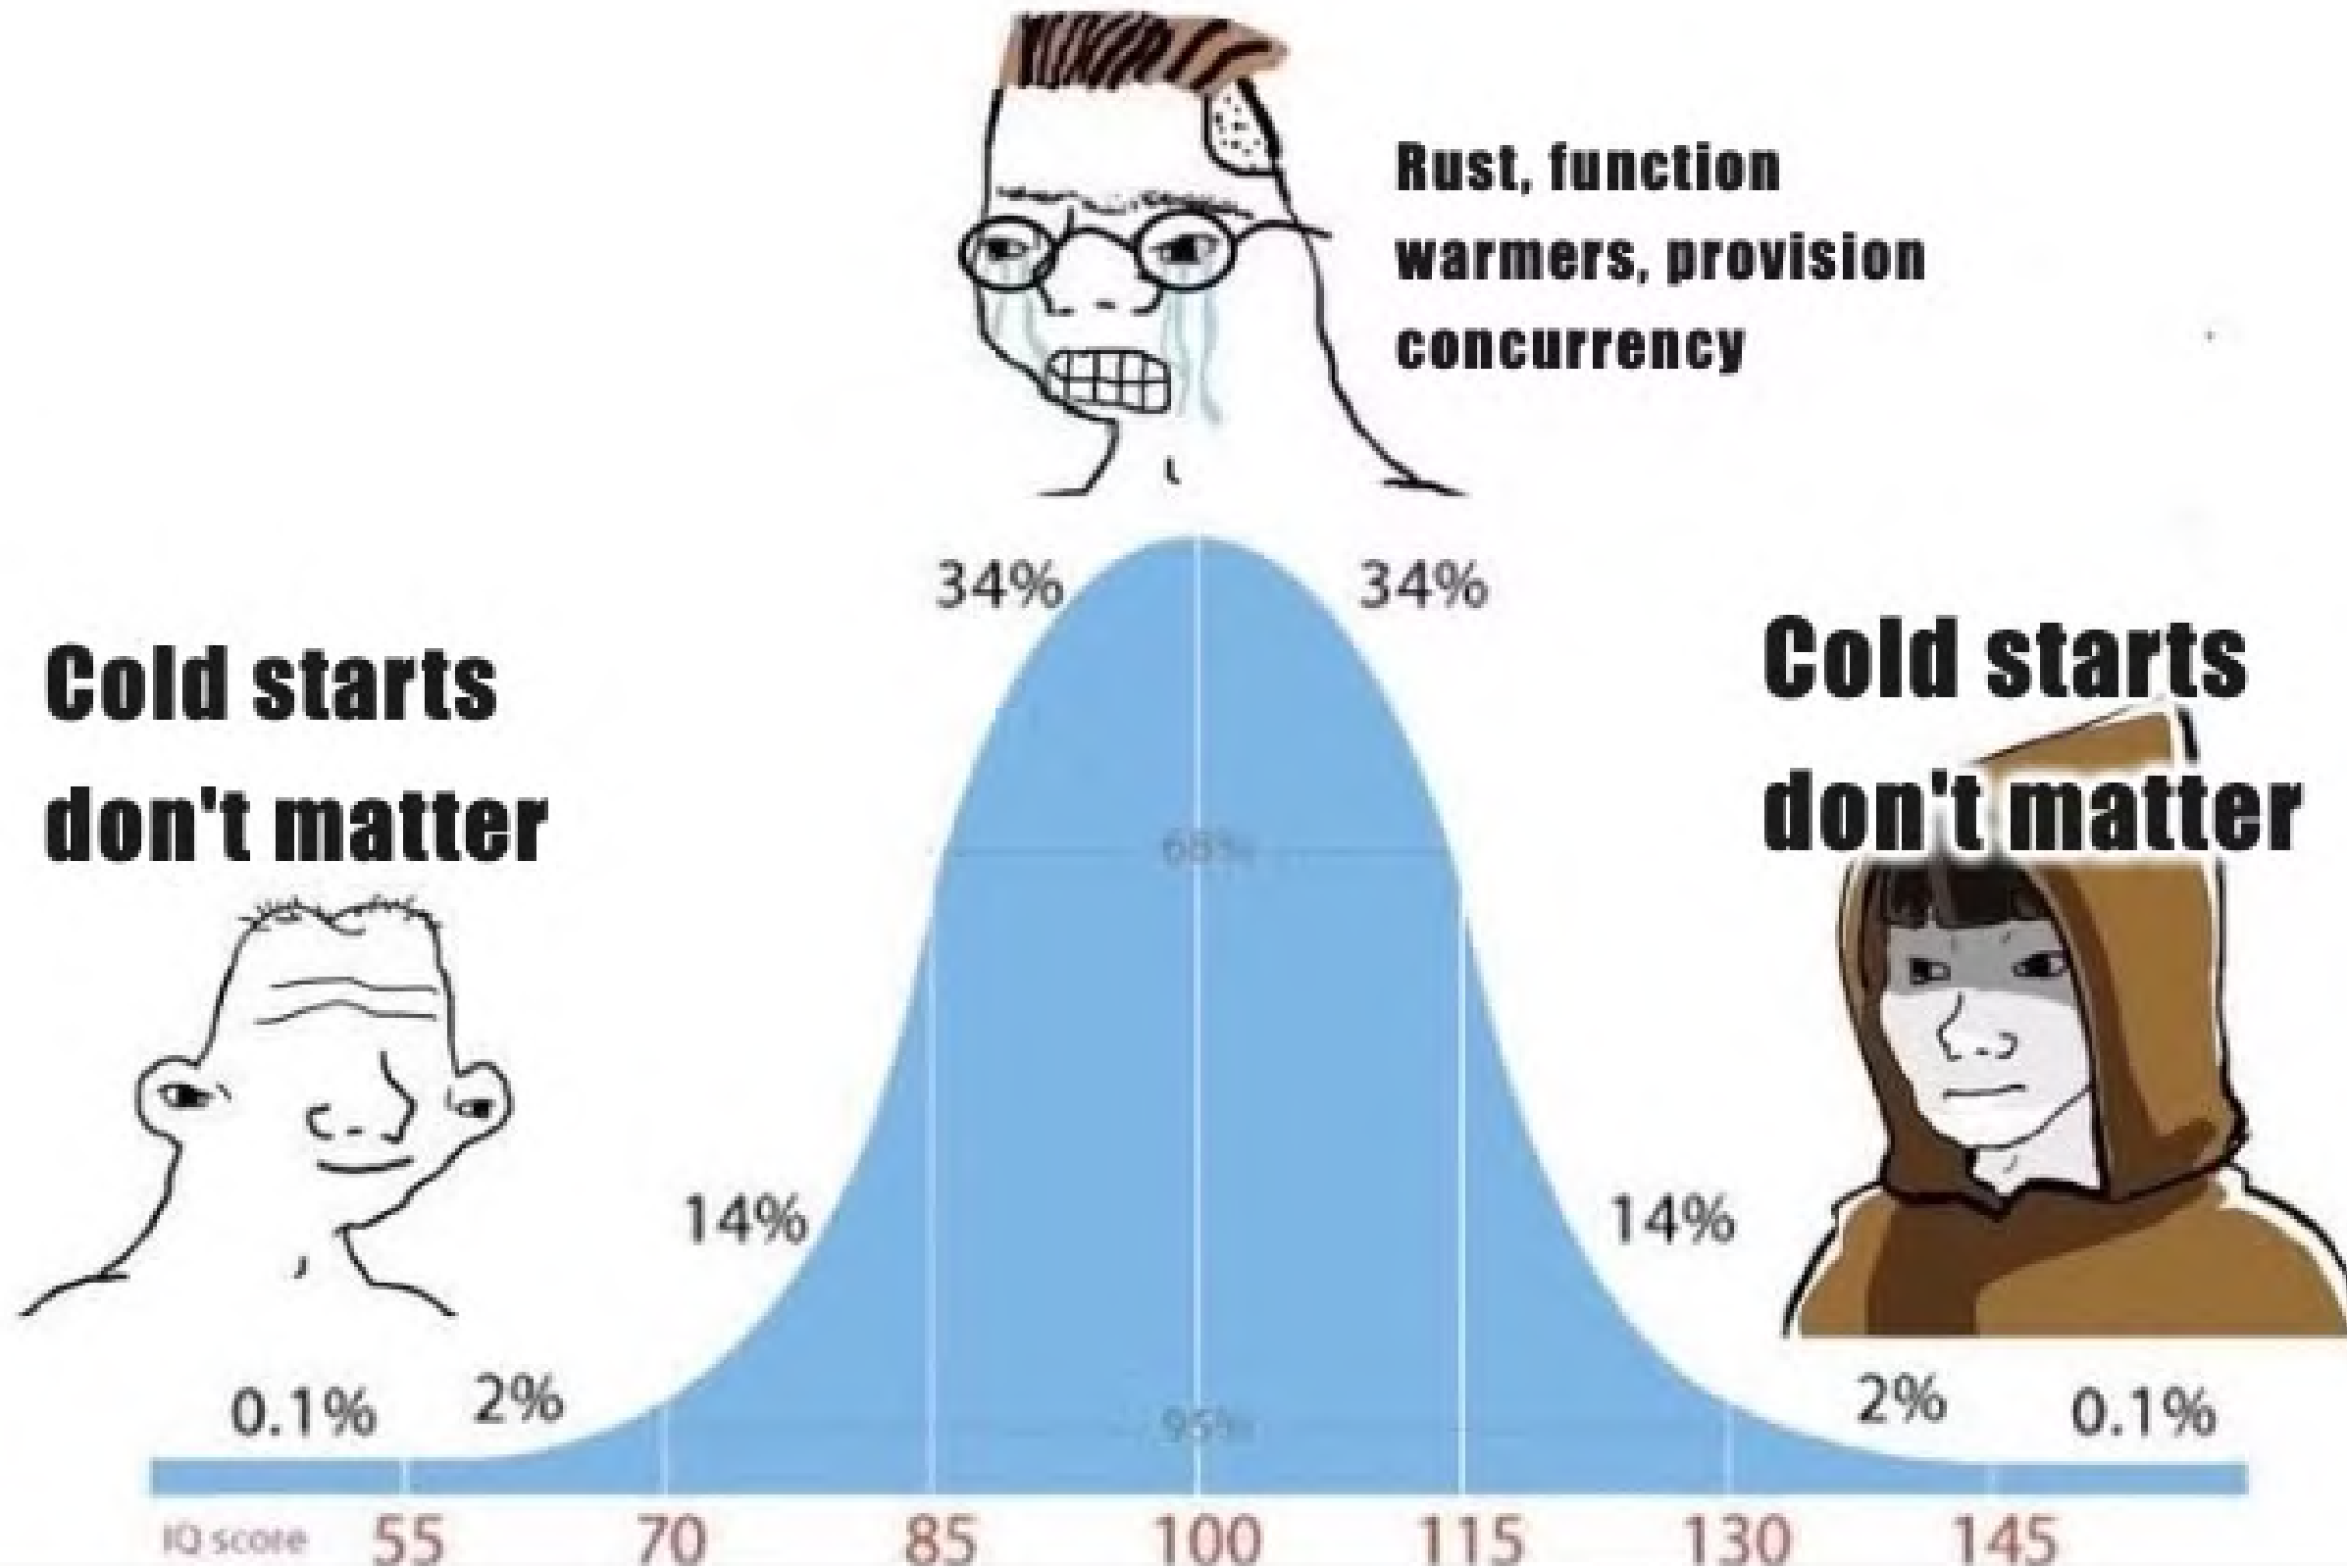
2024-06-18...

2024-06-17...

2024-06-16...

2024-06-15...

- Cold start
- But they c
- We use ca



P95Ms

479.9773

657.6846

722.8642

617.5303

657.6846

328.9003

527.5453



# Caching

WHY? | EDGE VS. NON-EDGE | VERCEL VS. API RESPONSE





# Caching

WHY?



**"Caching is a cheat code in distributed systems"**




YAN CUI



# Static Content

This one doesn't change often



visit blog post

blog.awsfundamentals.com/sfn

### Callback Patterns Let You Wait for User Input

Today, we're diving into the magic of the **callback pattern** with Step Functions, a feature that's like putting your workflow on pause. Think of it as a "Hold, please! 🙏" button for your tasks. The callback pattern lets a part of your workflow wait quietly until it's given the green light to go again.

```
graph LR; A[Get Data] --> B[Send to Admin]; B --> C[Wait for Callback]; C -- continue --> B; C --> D[Execute]
```

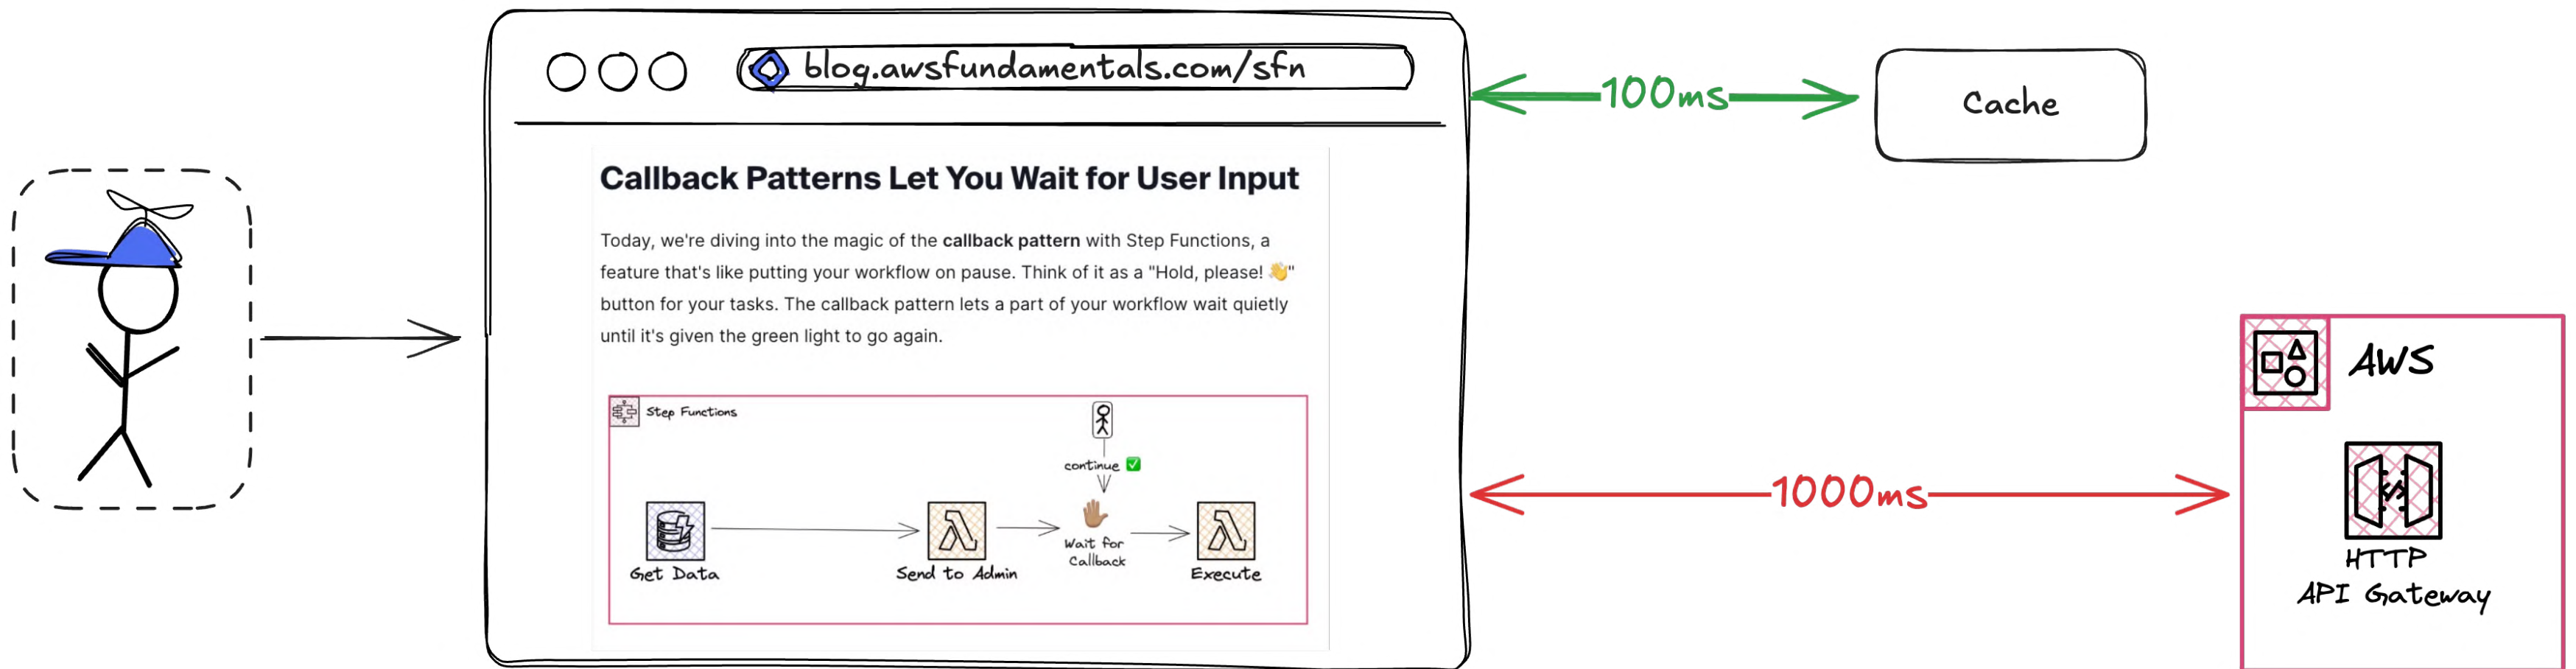
Cache

AWS

HTTP API Gateway




# Performance





# Protect Downstream Systems

*This is always the problem!*



blog.awsfundamentals.com/sfn

### Callback Patterns Let You Wait for User Input

Today, we're diving into the magic of the **callback pattern** with Step Functions, a feature that's like putting your workflow on pause. Think of it as a "Hold, please! 🙏" button for your tasks. The callback pattern lets a part of your workflow wait quietly until it's given the green light to go again.

```
graph LR; A[Get Data] --> B[Send to Admin]; B --> C[Wait for Callback]; C --> D[Execute];
```

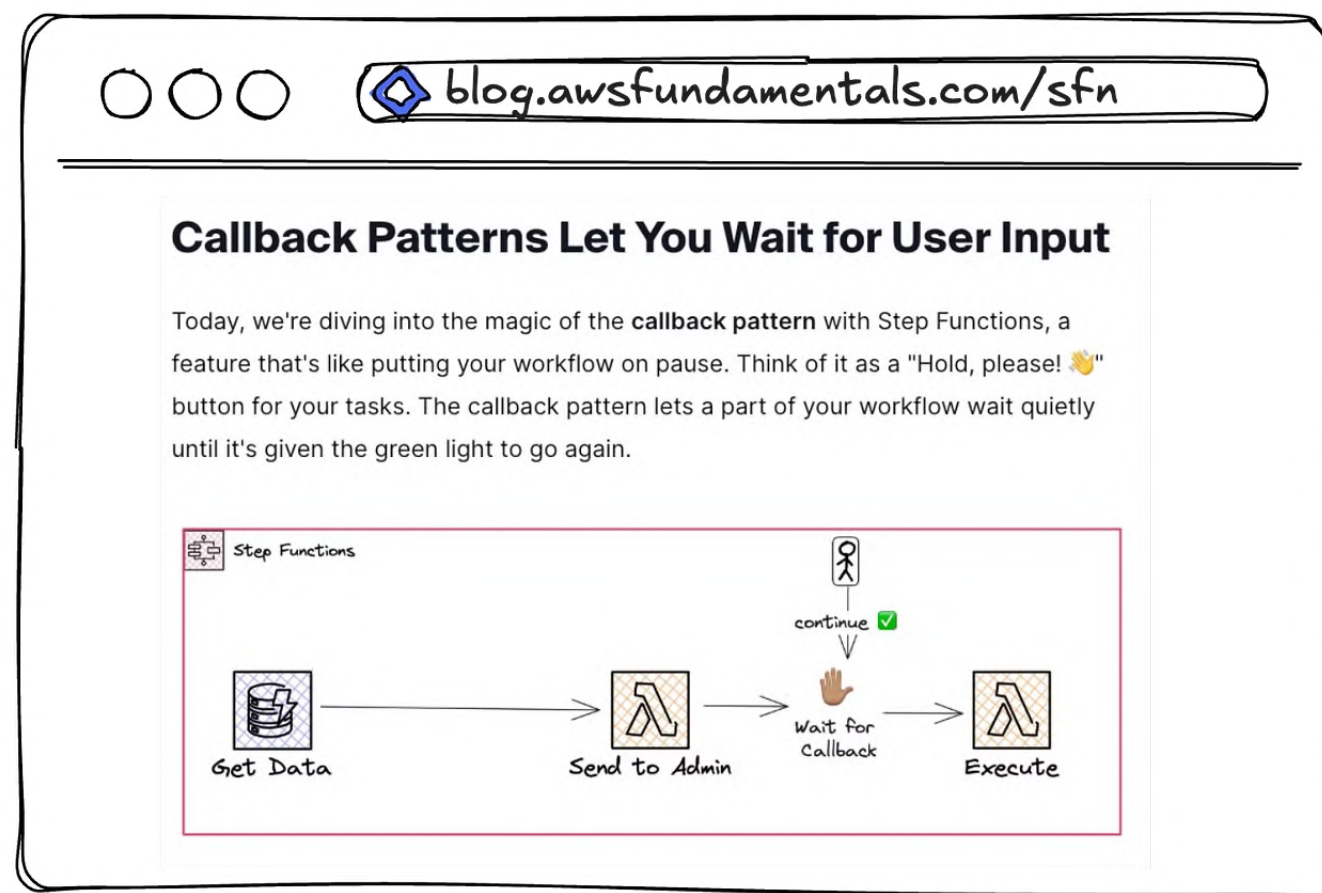

Cache

AWS  
HTTP API Gateway



MongoDB Atlas  
MongoDB



# Less Computation = More Money



Fewer requests 💰



# Why caching TLDR;

- static content doesn't change often
- performance
- protect downstream system
- less computation

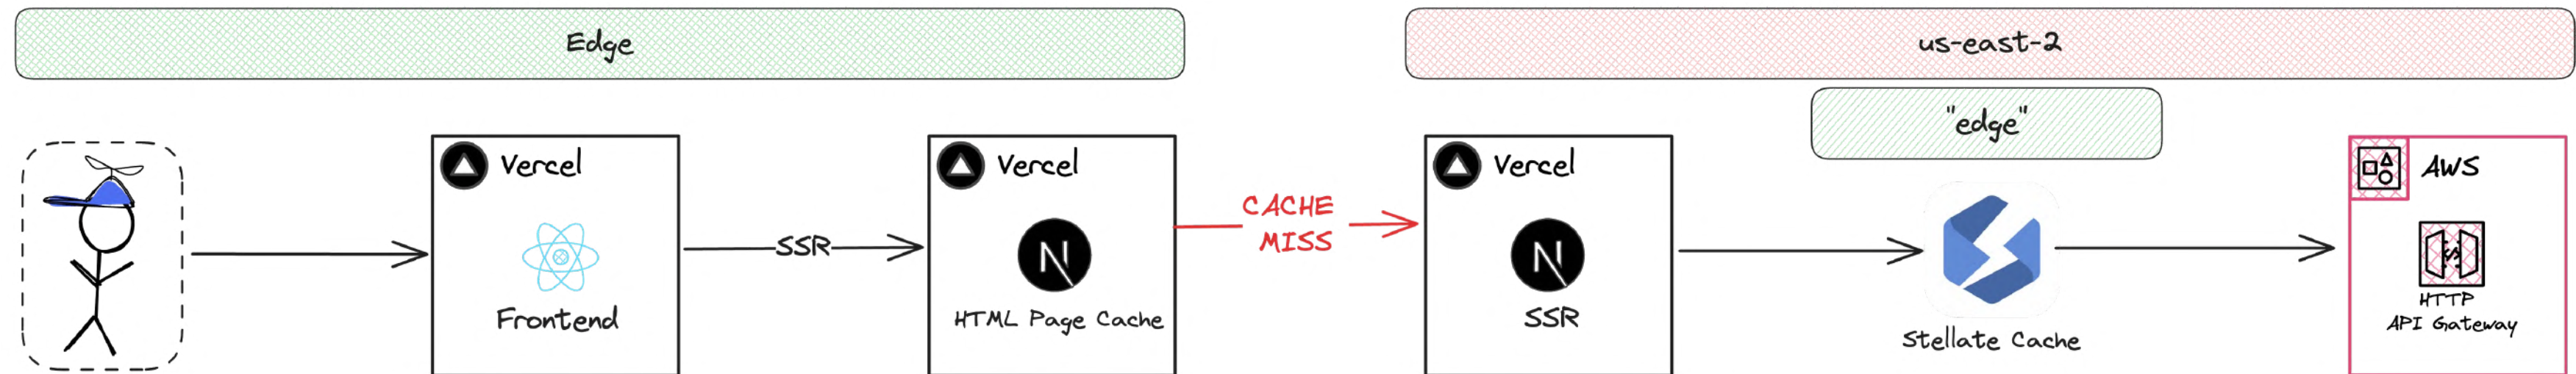


# Caching

EDGE VS. NON-EDGE

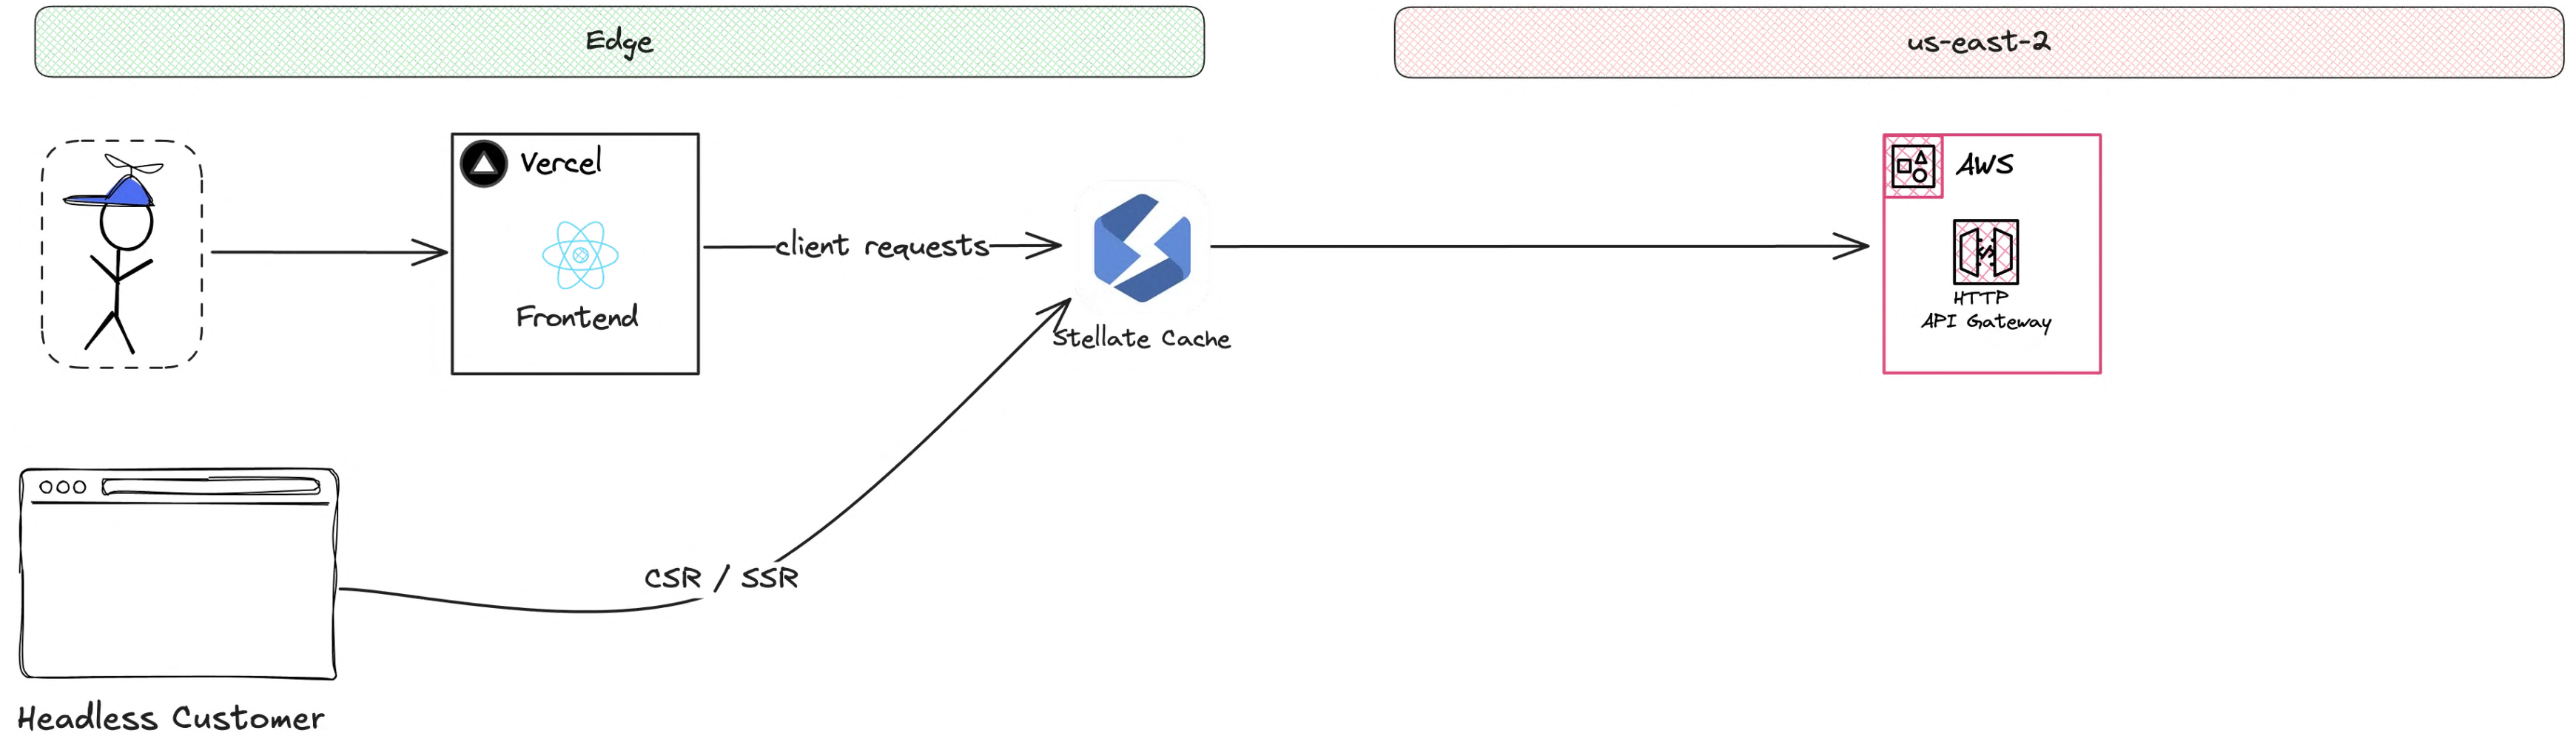


# Edge vs. non-edge - Server-Side Rendering



- still faster
- but not edge
- not 100% optimal
- more security than performance


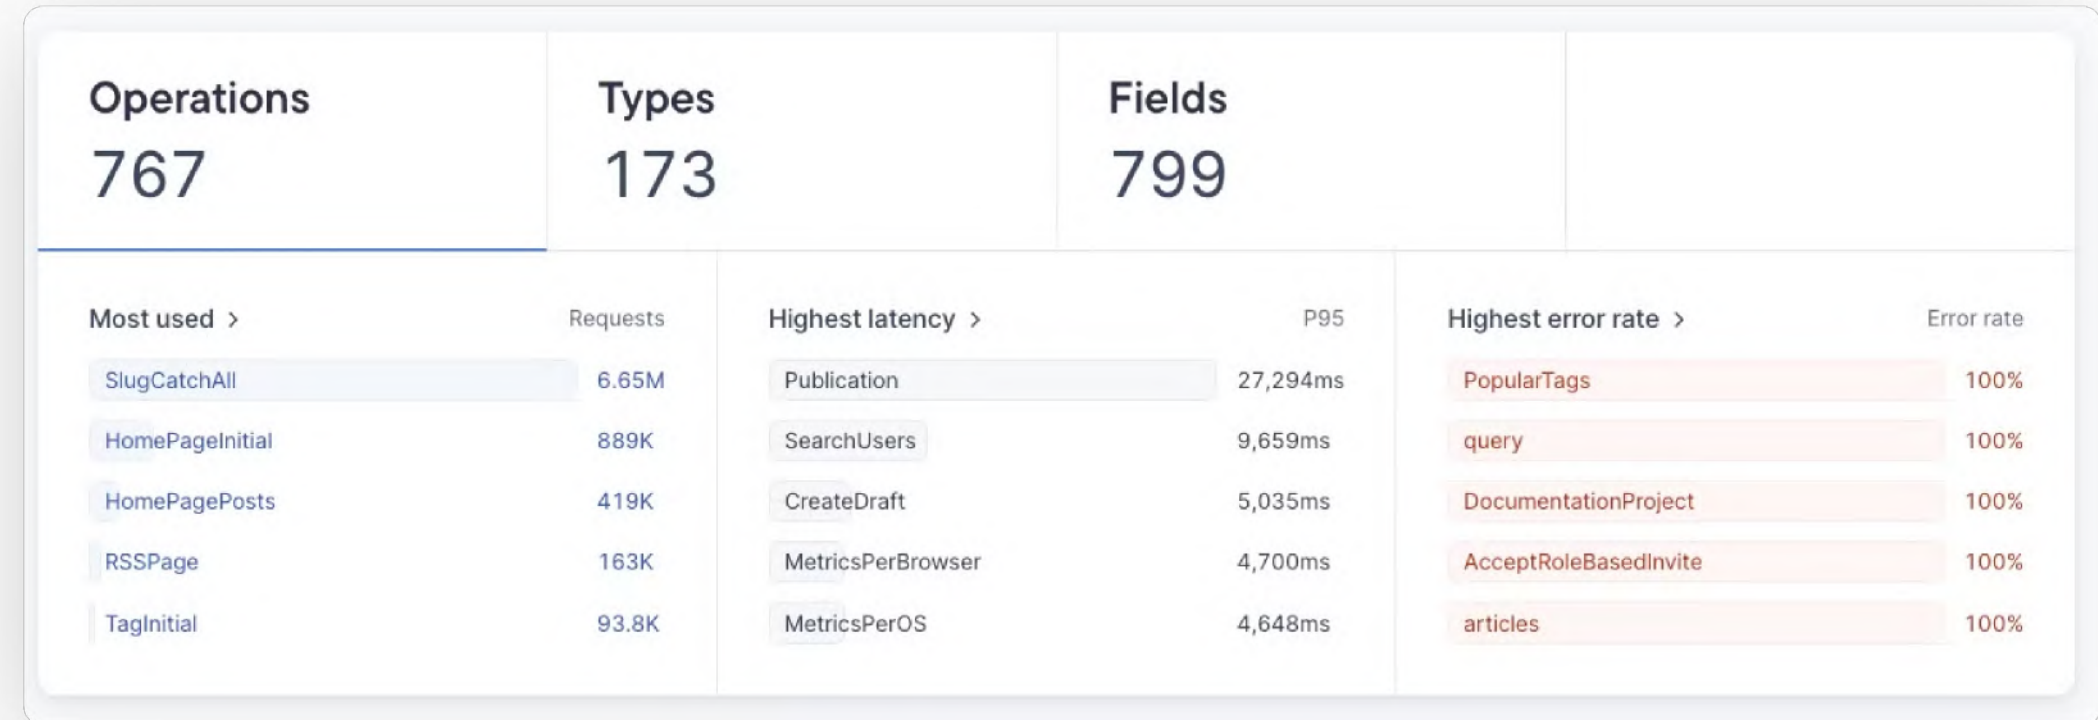
# Edge vs. non-edge - Client-Side Rendering | Headless



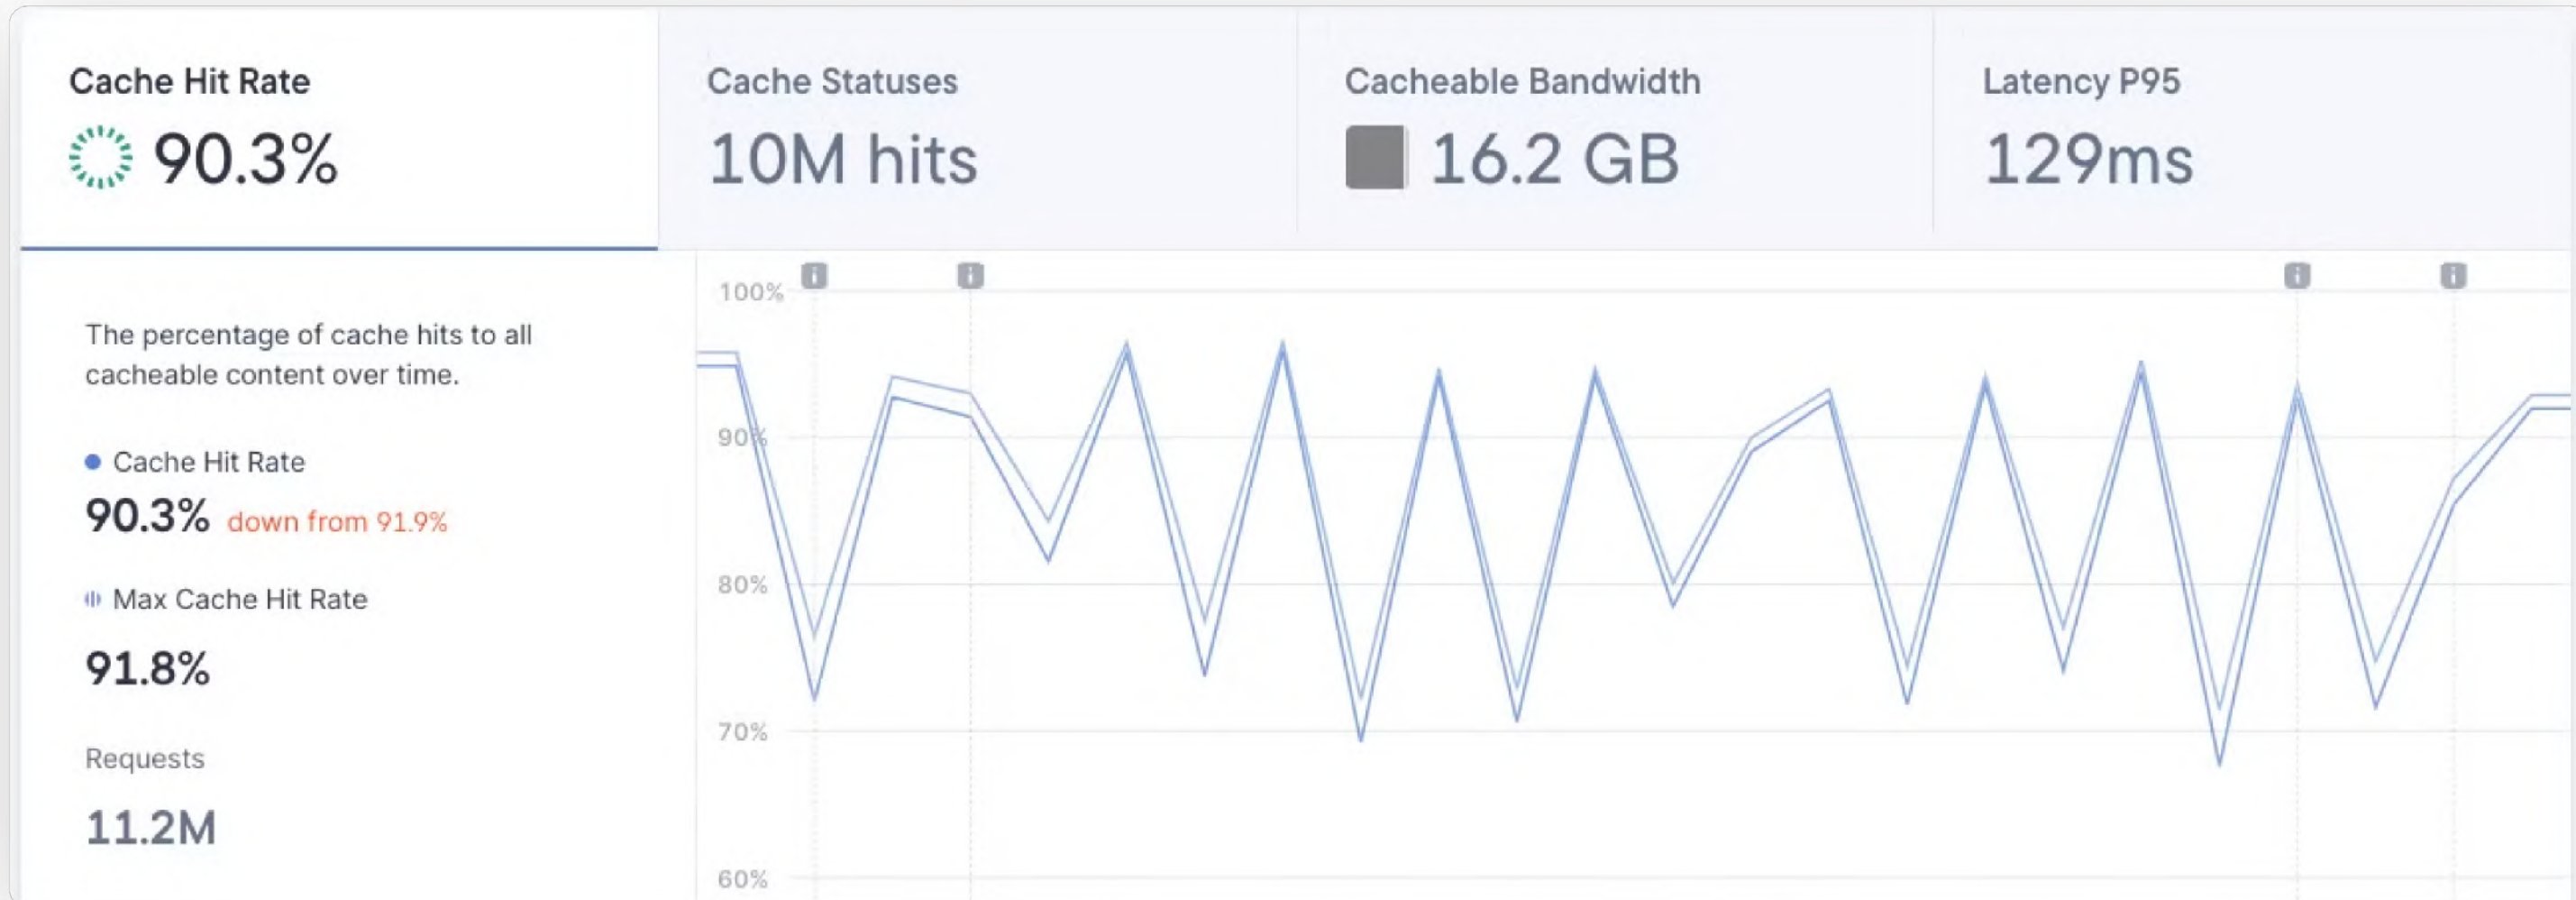
# Stellate

Stellate is a third-party provider that provides GraphQL Tooling:

- Caching
- Metrics
- Rate Limits



# Cache Hit Rate





# Lessons Learned

ERROR HANDLING | ABUSING | CACHING IS HARD | NAMING IS HARD

# GQL Lessons Learned from a non-GQL pro

## GQL Learnings:

- Error handling is hard
- Stand on the shoulder of giants (sorry Luc) → Have Design Guidelines
- Breaking changes → Mobile Apps
- Versioning
- n+1 Problem: Use Dataloaders
- You offer data → Frontend will use **all data** → Education

# General Lessons Learned from a non-GQL pro

## Public API Learnings

- Your API will be bombarded
- The issue is **always** the downstream system
- Polling is easier than realtime subscriptions
- Code will always be a mess → Constant improvements with every PR
- Caching is hard
- Naming is hard

# Resources

[You probably don't need GraphQL - Max Stoiber \(Stellate\)](#)

[200 OK! Error Handling in GraphQL](#)

[The Hidden Cost of GraphQL & NodeJS](#)

[I am done with GraphQL after 6 years](#)

[Introduction to Amazon API Gateway and AWS AppSync](#)





# Want to make a presentation like this one?

Start with a fully customizable template, create a beautiful deck in minutes, then easily share it with anyone.

[Create a presentation \(It's free\)](#)