APIs at **scale**

Matthew Clark

@matthew1000

BBC

Paul Caporn

@drpacman72

Designing for **frontends**

Designing for **change**

Designing for **scale**

Designing for **validation**

BBC Bitesize

BBC WEATHER

BBC SPORT

BBC NEWS

BBC NEWS РУССКАЯ СЛУЖБА

BBC Teach

BBC IDEAS

BBC

BBC MUNDO

BBC عربي

BBC cymru fyw

BBC FOOD

BBC MUSIC

BBC LEARNING ENGLISH

BBC News, the #1 most popular news site in the world

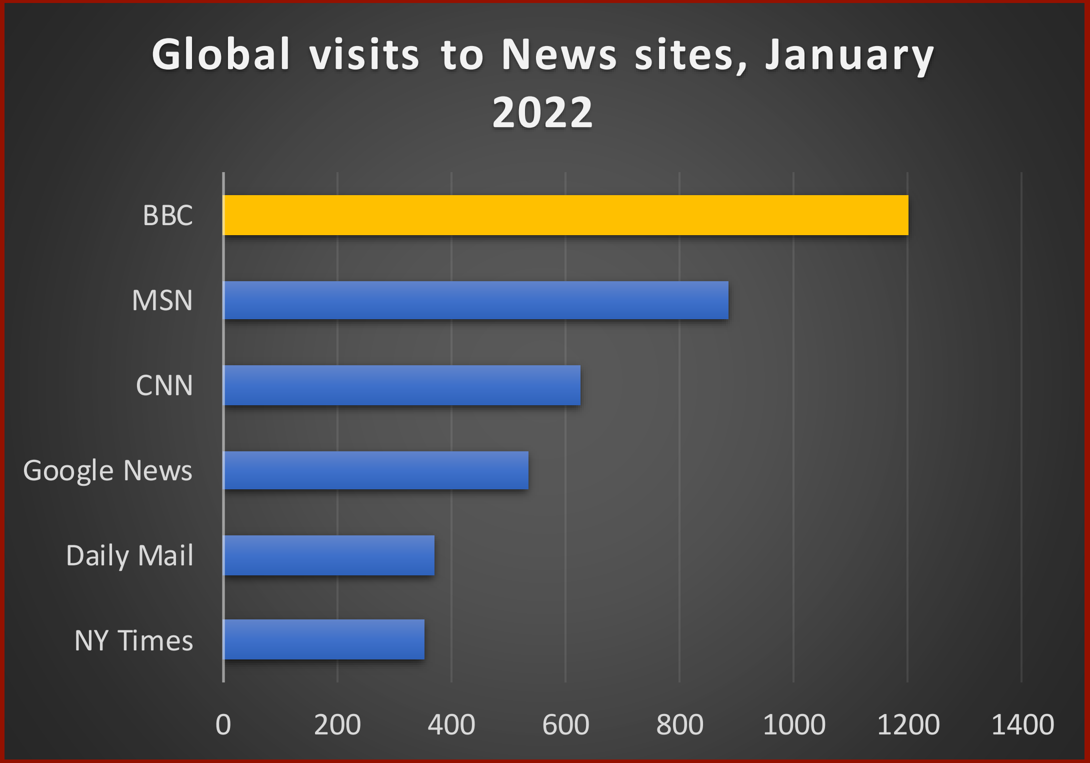

Source: Press Gazette / Similalrweb
https://tinyurl.com/most-popular-news

BBC iPlayer, the #2 most popular online video service in the UK

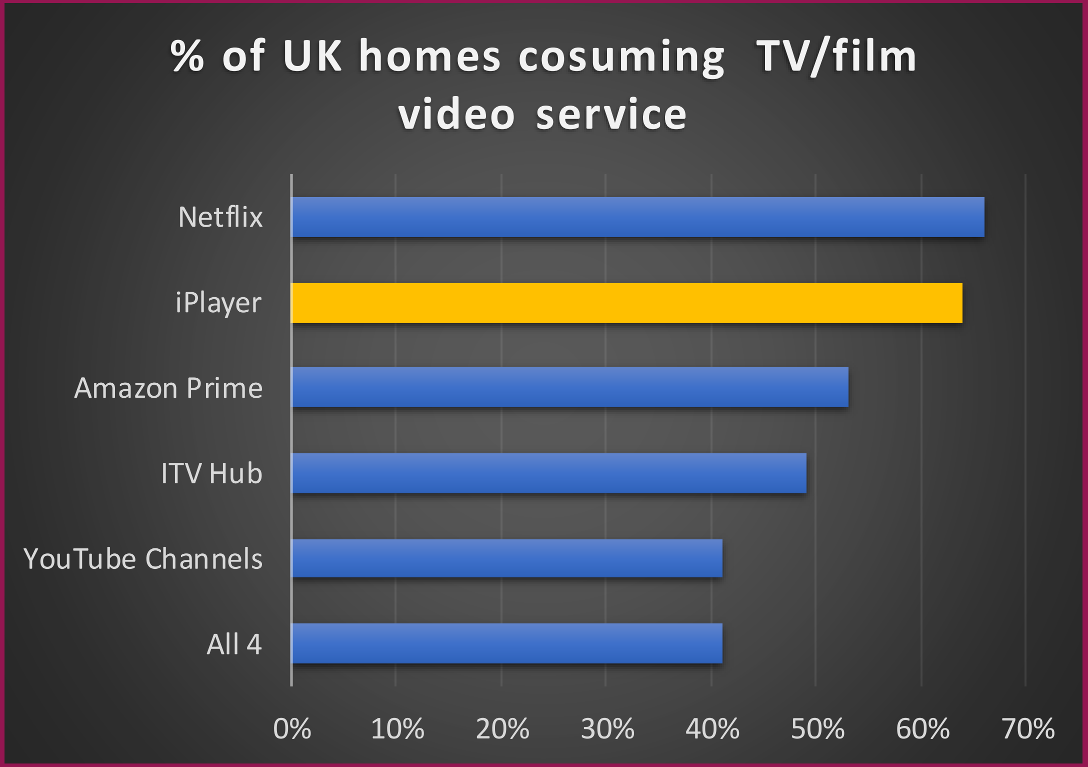

Source: Ofcom
https://tinyurl.com/ofcom-report-2021

Over 100 APIs power BBC Online

API name	primary response type	Compute type	directly powers audienc experience?	Publicly advertised?	Open to world?	Access control	Request interface type	Has event-driven?
App Business Layer- News	JSON	FABL	Yes	No	Yes	None	HTTP	No
App Business Layer - Sport	JSON	FABL	Yes	No	Yes	None	HTTP	No
Account Identity Command Service	JSON	EC2	Yes	No	No	Certificate	HTTP	No
Account token	JSON	EC2	Yes	No	No	Cookie	HTTP	No
APPW	JSON	S3	No	No	No	IAM	S3	Yes
Archive Services API	JSON	Serverless	No	Yes	Yes	Key	HTTP	No
Ares	JSON	EC2	No	No	No	Certificate	HTTP	No
Article domain	JSON	FABL	No	No	No	Certificate	HTTP	No
Audiences Portal API	JSON	EC2	No	No	No	Certificate	HTTP	No
BBC Ideas	JSON	Morph	Yes	No	No	Certificate	HTTP	Yes
BBC Live	JSON	Morph	Yes	No	Yes	None	HTTP	Yes
Belfrage	Multiple	EC2	Yes	No	No	None	HTTP	Yes
Binary Store	Multiple	EC2	Yes	No	Yes	None	HTTP	No
Blocks (like Arts)	JSON	Morph	Yes	No	No	Certificate	HTTP	Yes
Broadcast interactive data	JSON	EC2	No	No	Yes	Certificate	HTTP	Yes
Camino (News recs)	JSON	EC2	No	No	No	Certificate	HTTP	Yes
Comments	JSON		Yes	No	Yes	Token	HTTP	No
Content Summaries	JSON	FABL	No	No	No	Certificate	HTTP	No
Cosmos API	JSON	EC2	No	No	No	Certificate	HTTP	No
Counting Service	JSON	Serverless	Yes	No	Yes	None	HTTP	Yes

How many are serverless?

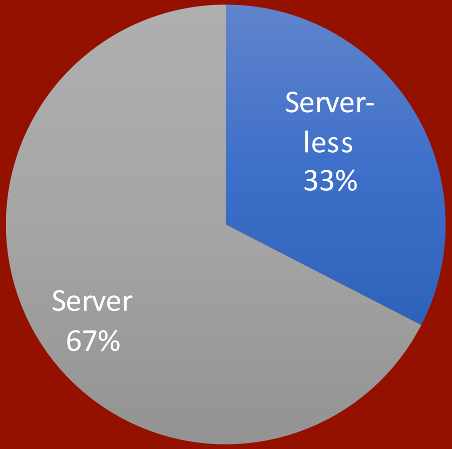

What means of access control?

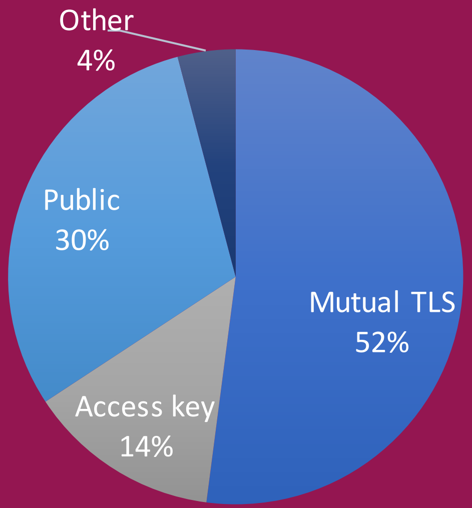

Event or request-driven?

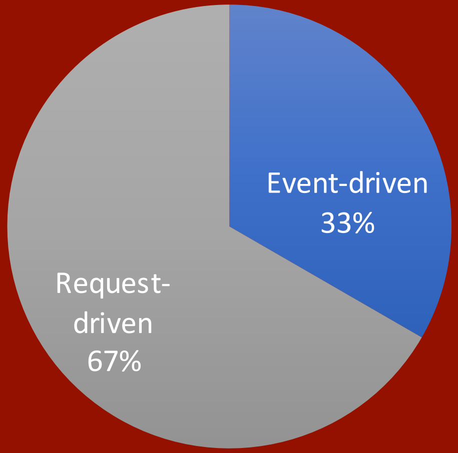

What response format?

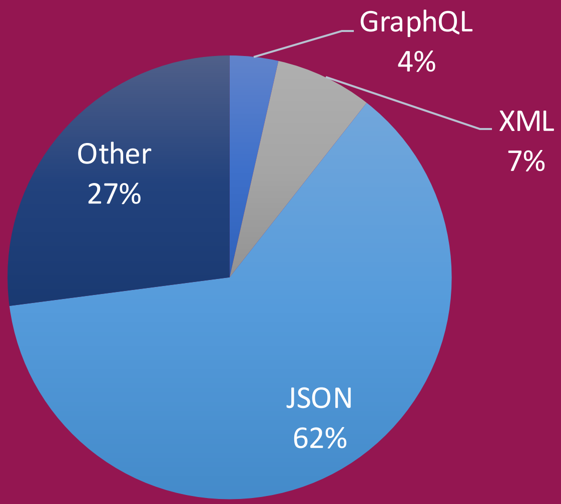

Designing for **frontends**

Micro-frontend API for TV’s

TV eco-system is fragmented (long tail)

Code using the API had sprawled…

Monolithic, spaghetti code

Micro-frontend API for TV’s

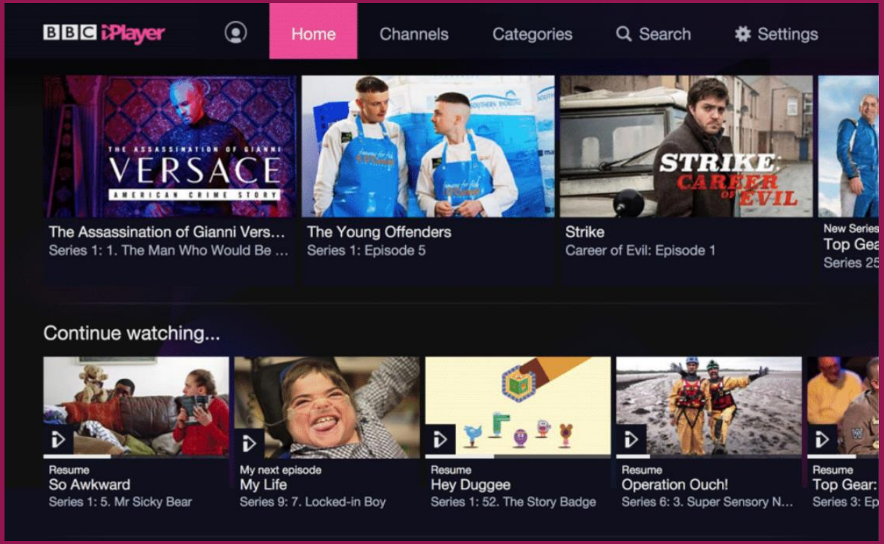

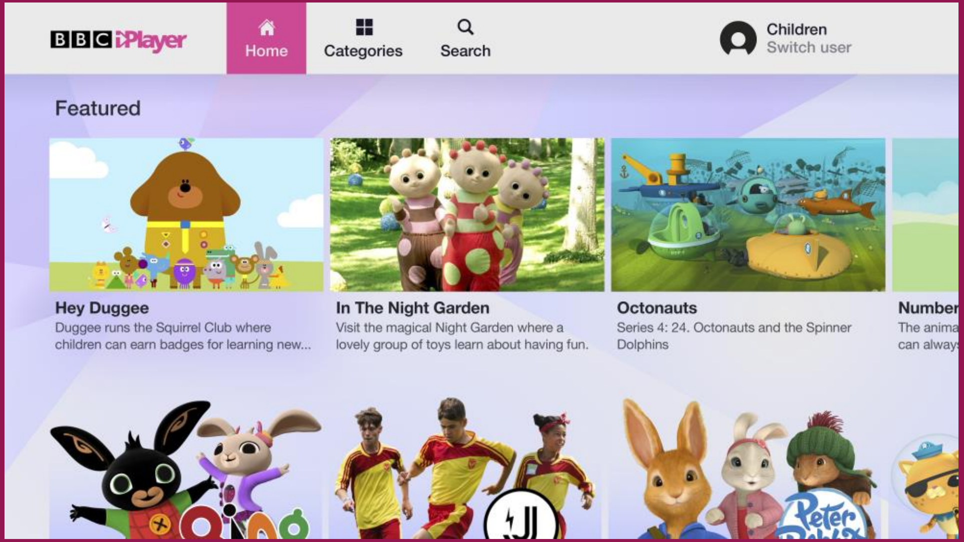
BBC iPlayer
Home
Categories
Search
Children
Switch user
Featured
Hey Duggee
Duggee runs the Squirrel Club where children can earn badges for learning new...
In The Night Garden
Visit the magical Night Garden where a lovely group of toys learn about having fun.
Octonauts
Series 4: 24. Octonauts and the Spinner Dolphins
Number
The anima
can alway

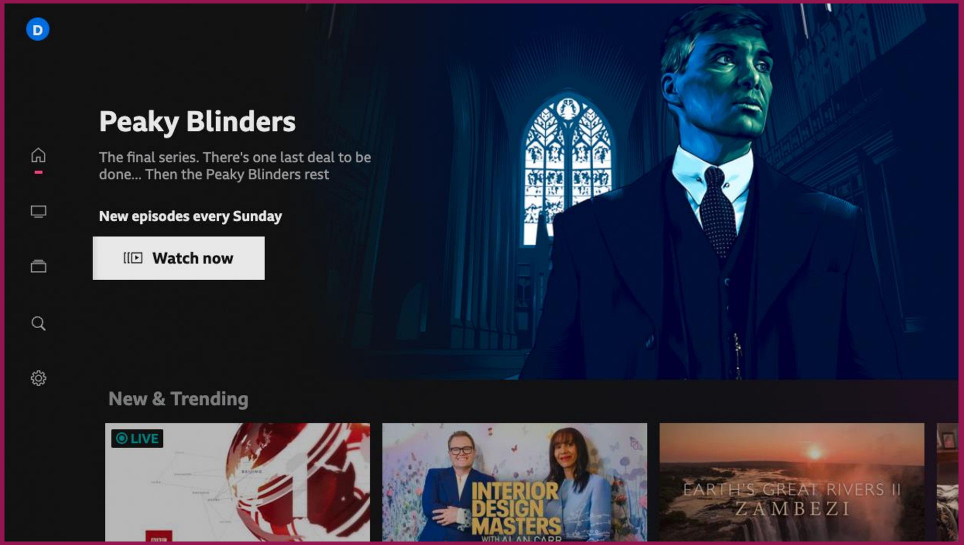
D
Peaky Blinders
The final series. There's one last deal to be done... Then the Peaky Blinders rest
New episodes every Sunday
Watch now
New & Trending
LIVE
INTERIOR DESIGN MASTERS
EARTH'S GREAT RIVERS II
ZAMBEZI

- Launch a kids app!
- Renew and iterate on the UI!
- But unclear boundaries / lack of a clear API

Micro-frontend API for TV’s

“One Love” project – enabling decoupled development

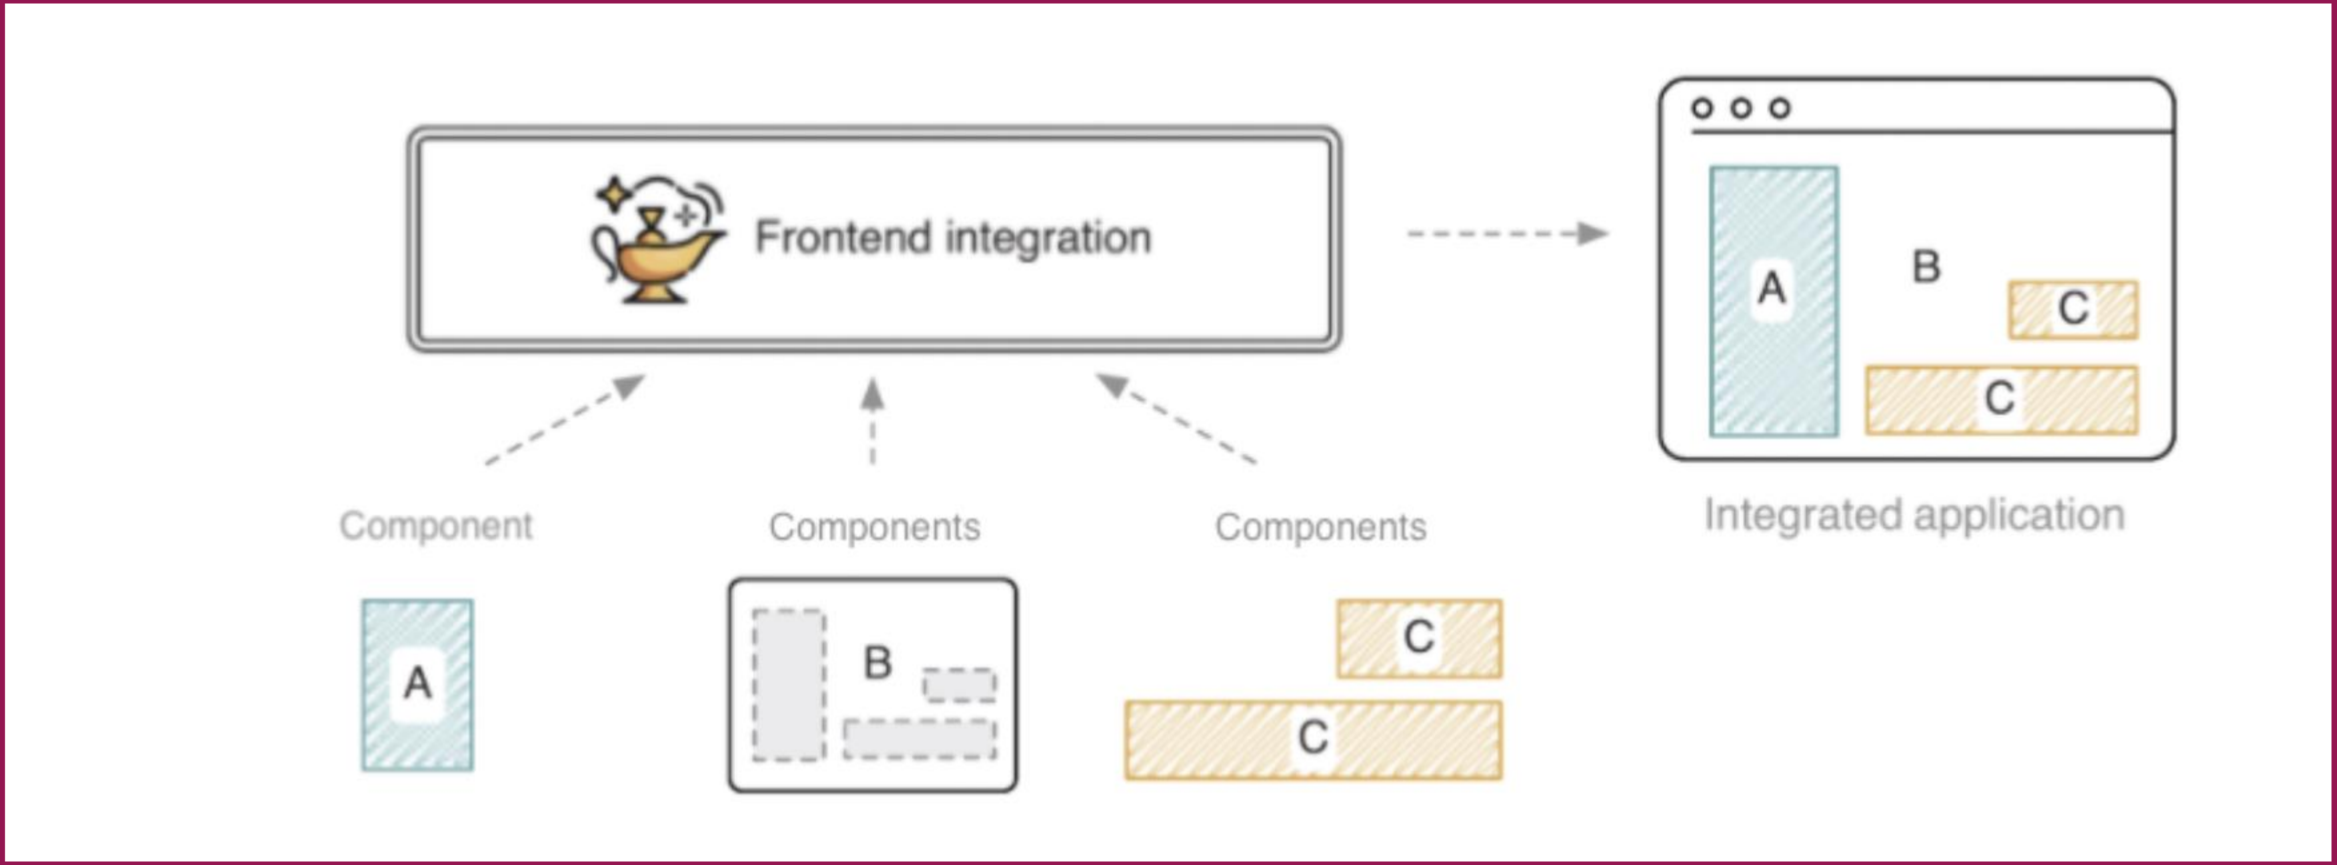

Micro-frontend API for TV’s

Component API – A small surface, clear separation of concerns

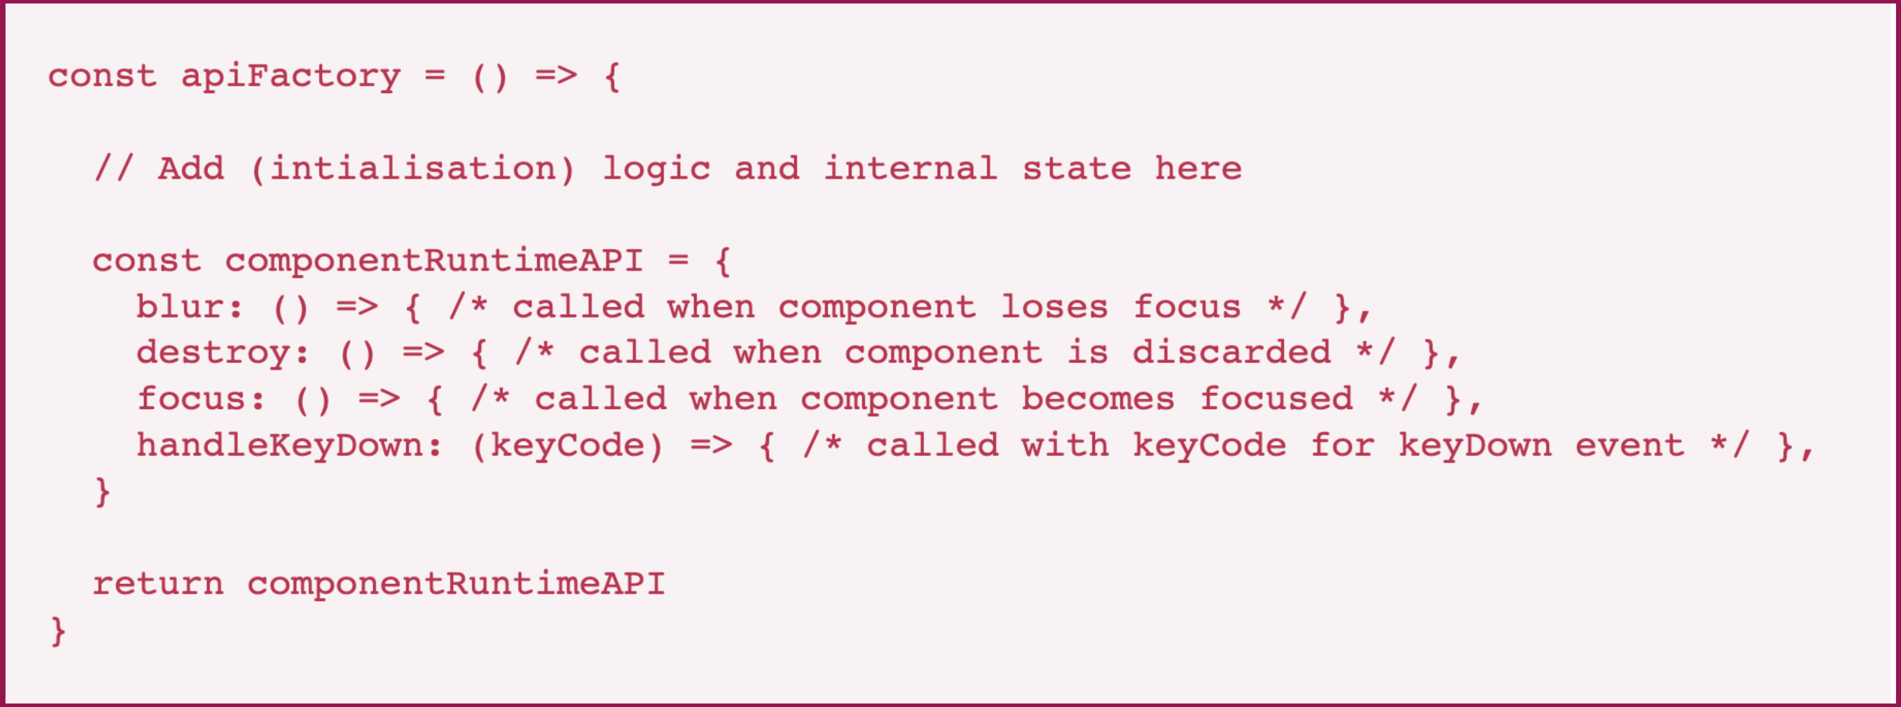

```
const apiFactory = () => {

  // Add (intialisation) logic and internal state here

  const componentRuntimeAPI = {
    blur: () => { /* called when component loses focus */ },
    destroy: () => { /* called when component is discarded */ },
    focus: () => { /* called when component becomes focused */ },
    handleKeyDown: (keyCode) => { /* called with keyCode for keyDown event */ },
  }

  return componentRuntimeAPI
}
```

Micro-frontend API for TV’s

Enabler for development of rich, data driven UI components in isolation

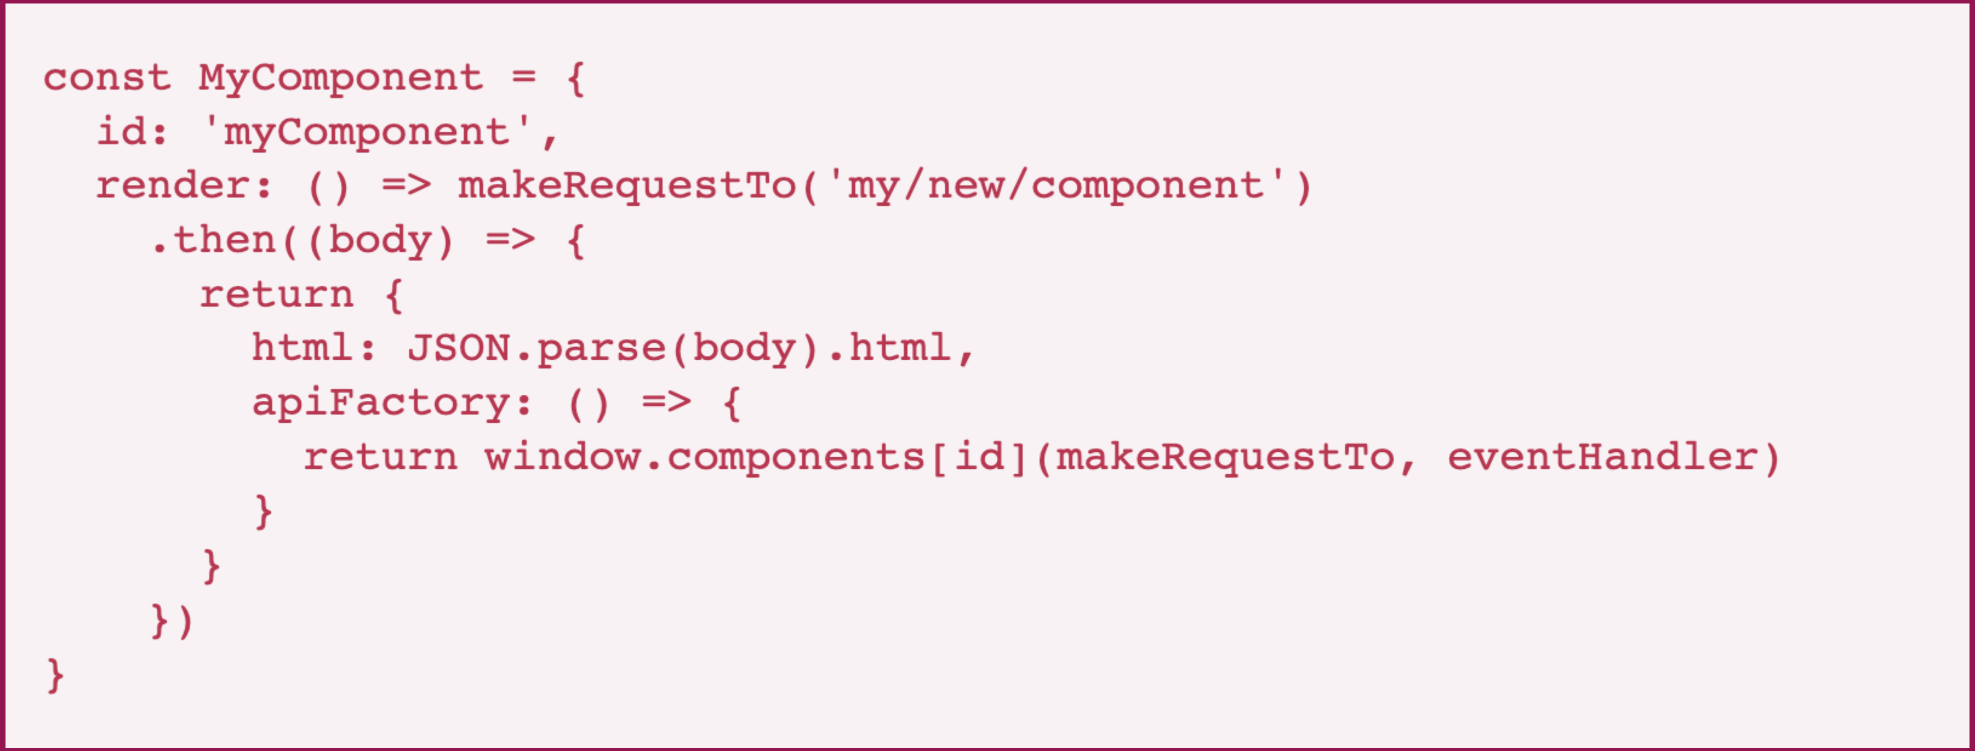

```
const MyComponent = {
  id: 'myComponent',
  render: () => makeRequestTo('my/new/component')
    .then((body) => {
      return {
        html: JSON.parse(body).html,
        apiFactory: () => {
          return window.components[id](makeRequestTo, eventHandler)
        }
      }
    })
}
```

Micro-frontend API for TV’s

Enabler for development of rich, data driven UI components in isolation

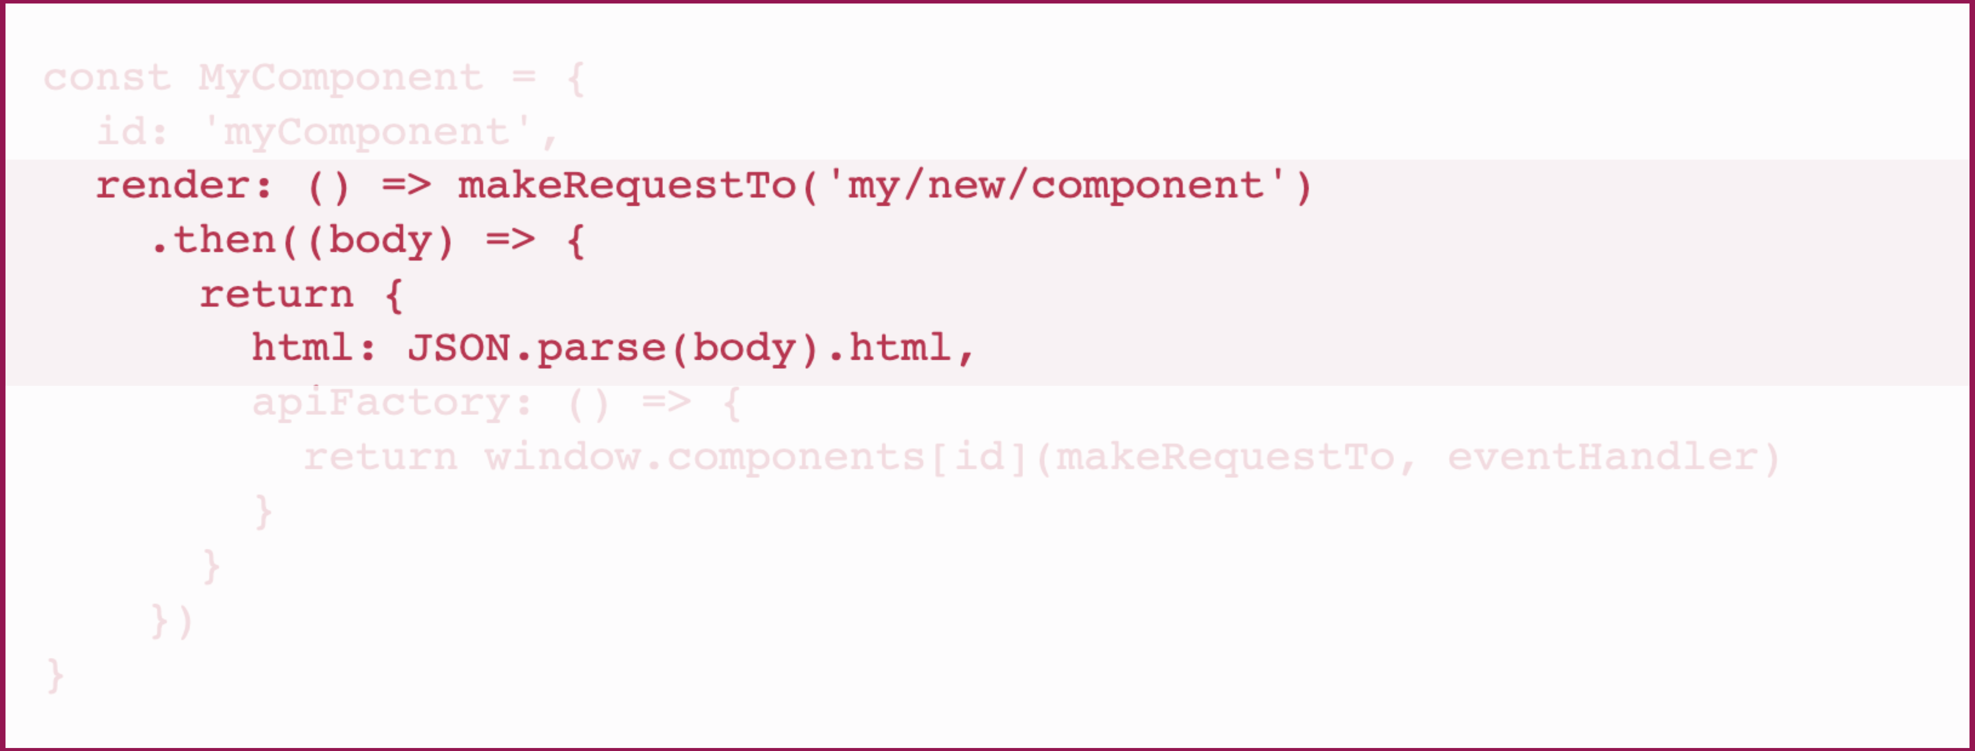

```
const MyComponent = {
  id: 'myComponent',
  render: () => makeRequestTo('my/new/component')
    .then((body) => {
      return {
        html: JSON.parse(body).html,
        apiFactory: () => {
          return window.components[id](makeRequestTo, eventHandler)
        }
      }
    })
}
```

Micro-frontend API for TV’s

Enabler for development of rich, data driven UI components in isolation

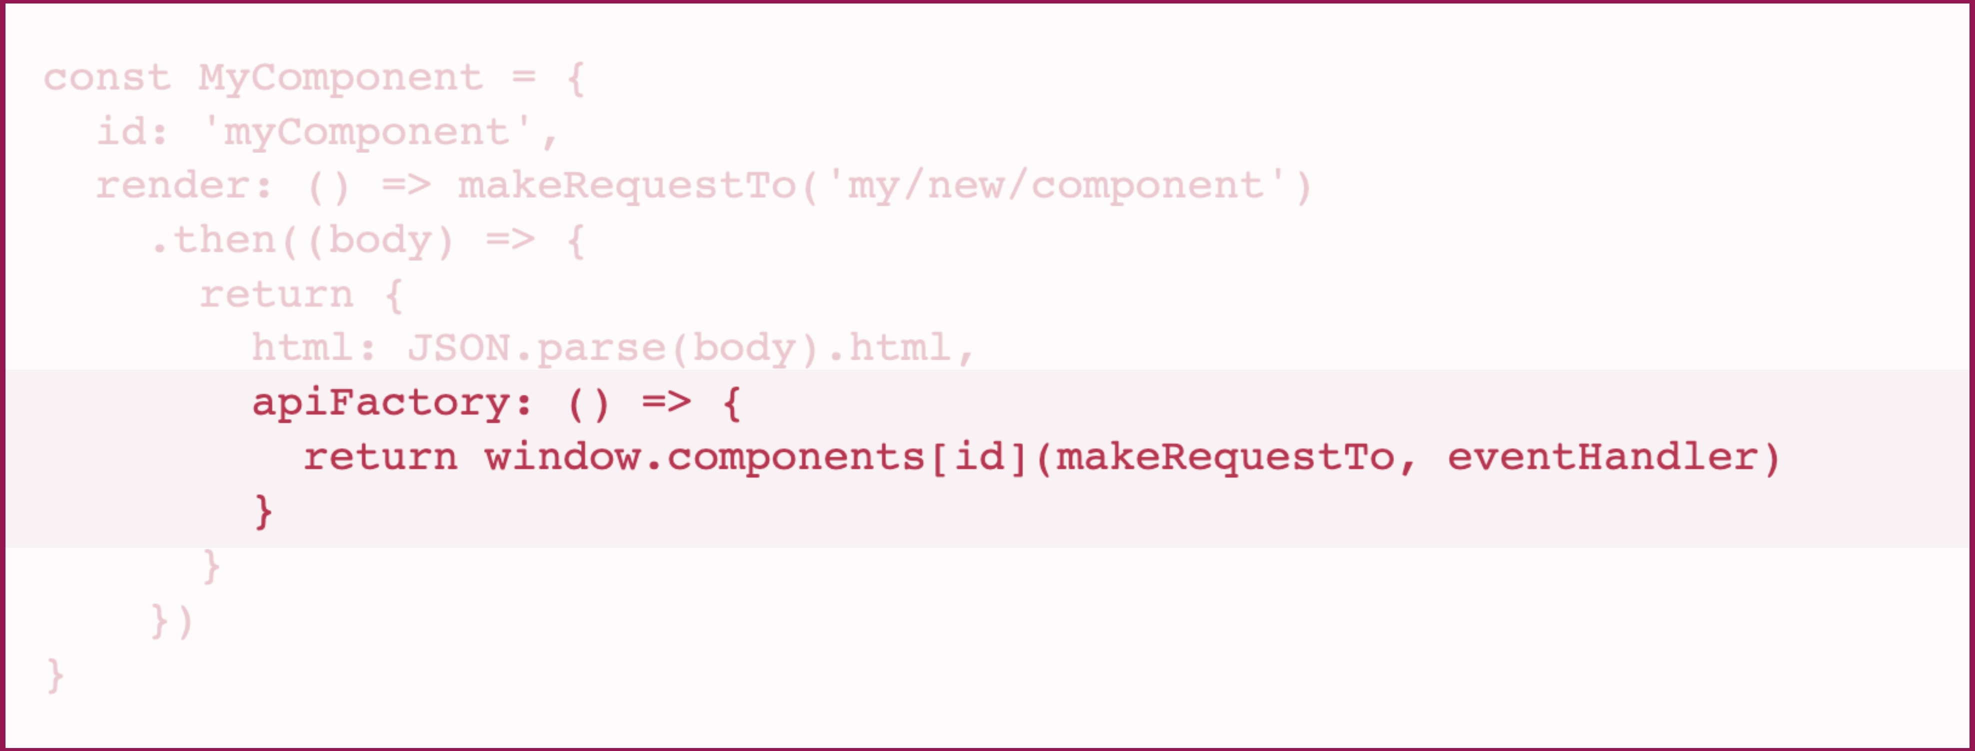

```
const MyComponent = {
  id: 'myComponent',
  render: () => makeRequestTo('my/new/component')
    .then((body) => {
      return {
        html: JSON.parse(body).html,
        apiFactory: () => {
          return window.components[id](makeRequestTo, eventHandler)
        }
      }
    })
}
```

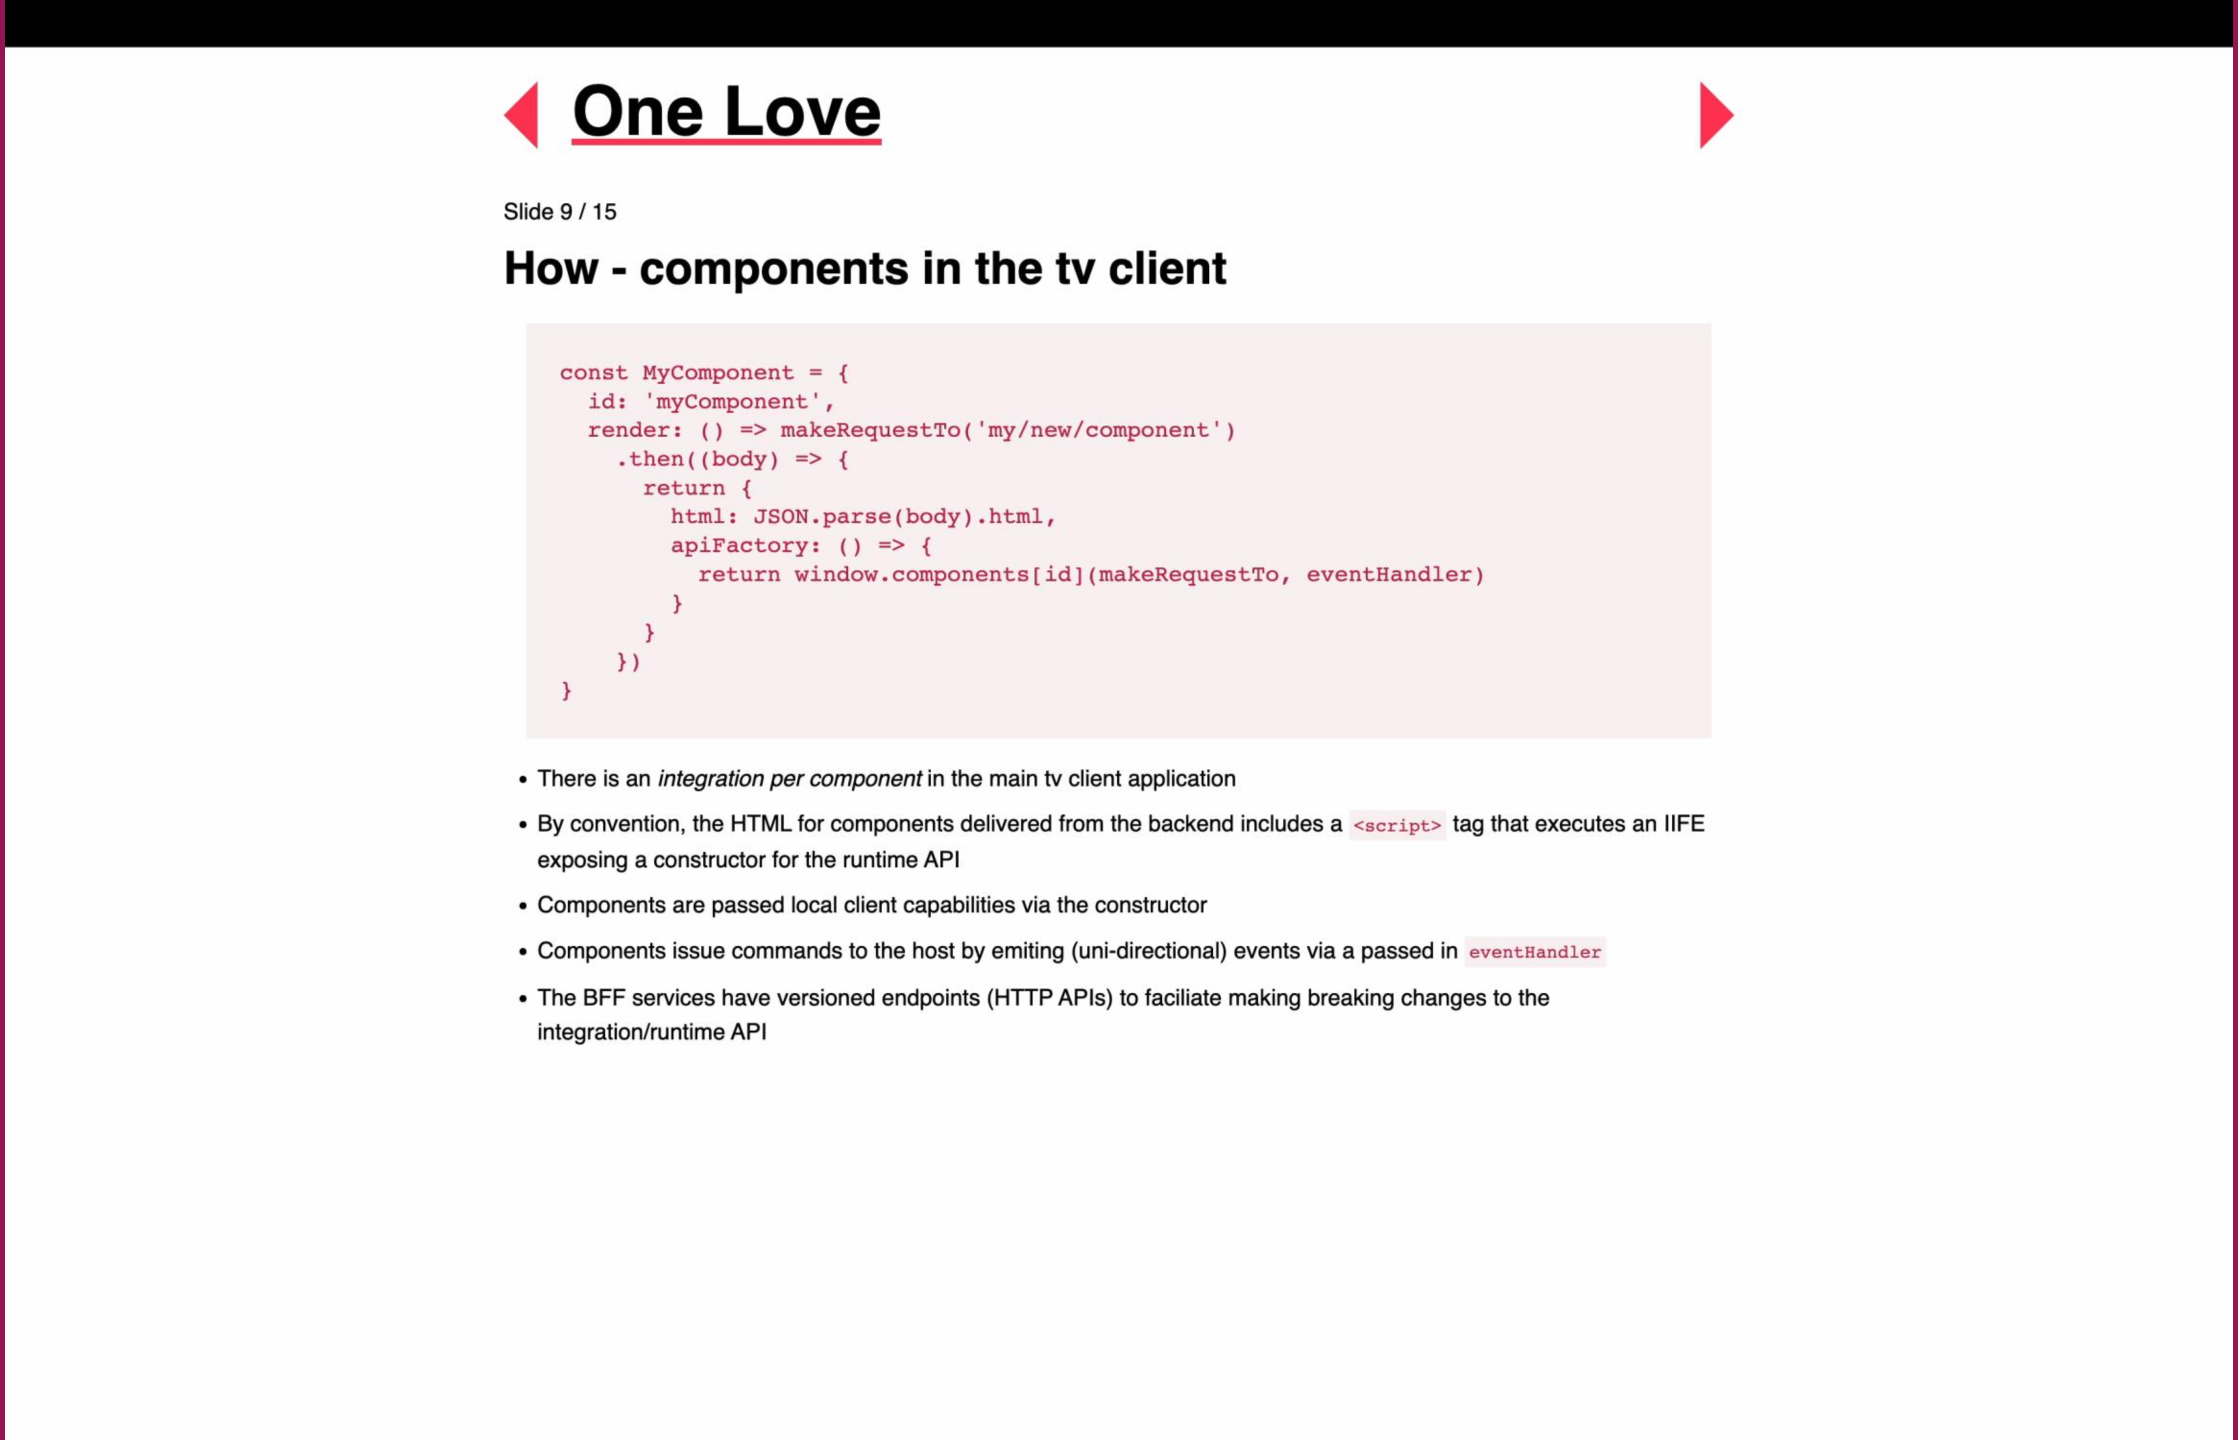

One Love

Slide 9 / 15

How - components in the tv client

```
const MyComponent = {
  id: 'myComponent',
  render: () => makeRequestTo('my/new/component')
    .then((body) => {
      return {
        html: JSON.parse(body).html,
        apiFactory: () => {
          return window.components[id](makeRequestTo, eventHandler)
        }
      }
    })
}
```

- There is an *integration per component* in the main tv client application
- By convention, the HTML for components delivered from the backend includes a `<script>` tag that executes an IIFE exposing a constructor for the runtime API
- Components are passed local client capabilities via the constructor
- Components issue commands to the host by emiting (uni-directional) events via a passed in `eventHandler`
- The BFF services have versioned endpoints (HTTP APIs) to faciliate making breaking changes to the integration/runtime API

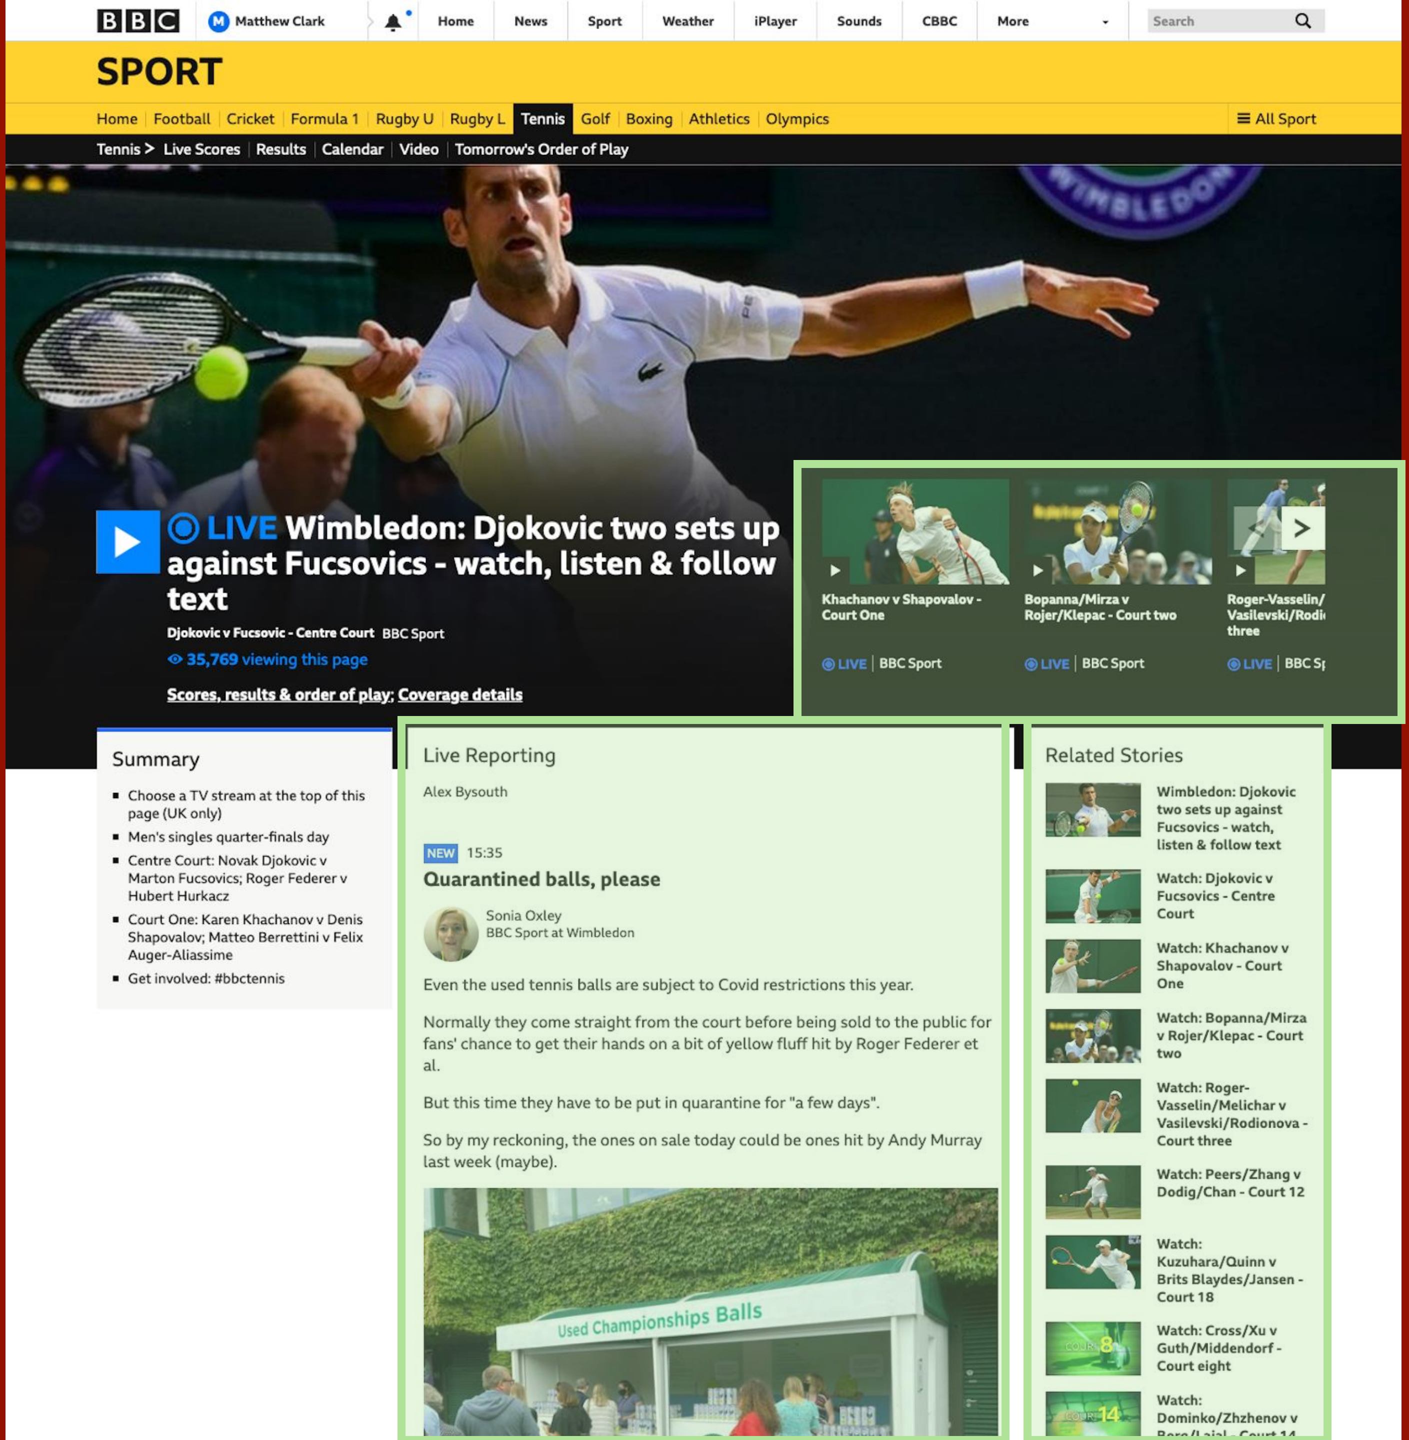

LIVE Wimbledon: Djokovic two sets up against Fucsovics - watch, listen & follow text

Djokovic v Fucsovic - Centre Court BBC Sport

35,769 viewing this page

Scores, results & order of play; Coverage details

Khachanov v Shapovalov - Court One
LIVE | BBC Sport

Bopanna/Mirza v Rojer/Klepac - Court two
LIVE | BBC Sport

Roger-Vasselin/ Vasilevski/Rodi three
LIVE | BBC Sp

Summary

- Choose a TV stream at the top of this page (UK only)
- Men's singles quarter-finals day
- Centre Court: Novak Djokovic v Marton Fucsovics; Roger Federer v Hubert Hurkacz
- Court One: Karen Khachanov v Denis Shapovalov; Matteo Berrettini v Felix Auger-Aliassime
- Get involved: #bbctennis

Live Reporting

Alex Bysouth

NEW 15:35

Quarantined balls, please

Sonia Oxley
BBC Sport at Wimbledon

Even the used tennis balls are subject to Covid restrictions this year.

Normally they come straight from the court before being sold to the public for fans' chance to get their hands on a bit of yellow fluff hit by Roger Federer et al.

But this time they have to be put in quarantine for "a few days".

So by my reckoning, the ones on sale today could be ones hit by Andy Murray last week (maybe).

Related Stories

- Wimbledon: Djokovic two sets up against Fucsovics - watch, listen & follow text
- Watch: Djokovic v Fucsovics - Centre Court
- Watch: Khachanov v Shapovalov - Court One
- Watch: Bopanna/Mirza v Rojer/Klepac - Court two
- Watch: Roger-Vasselin/Melichar v Vasilevski/Rodionova - Court three
- Watch: Peers/Zhang v Dodig/Chan - Court 12
- Watch: Kuzuhara/Quinn v Brits Blaydes/Jansen - Court 18
- Watch: Cross/Xu v Guth/Middendorf - Court eight
- Watch: Dominko/Zhzhenov v

APIs doing one thing well

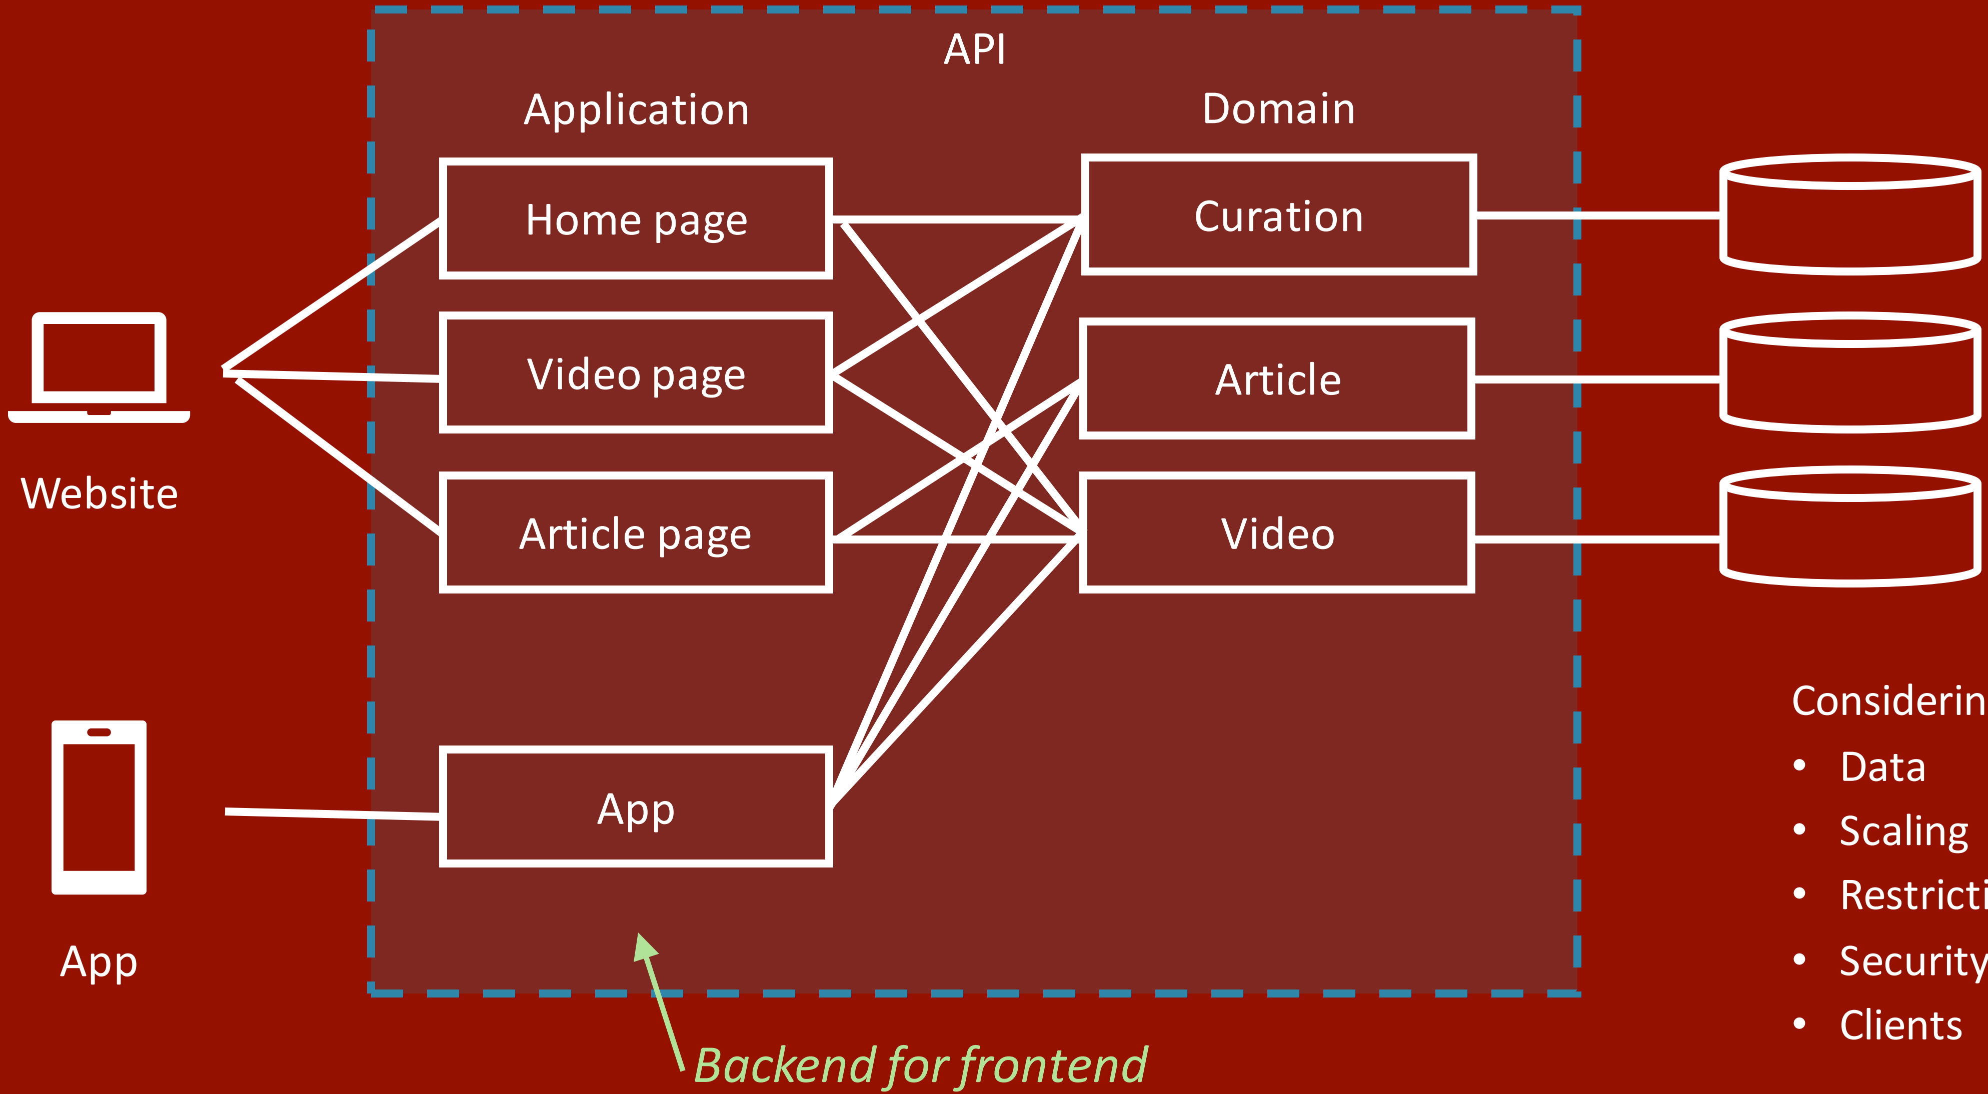

Backend for frontend

Considering the…

- Data
- Scaling
- Restrictions
- Security
- Clients

Thanks, X-ray

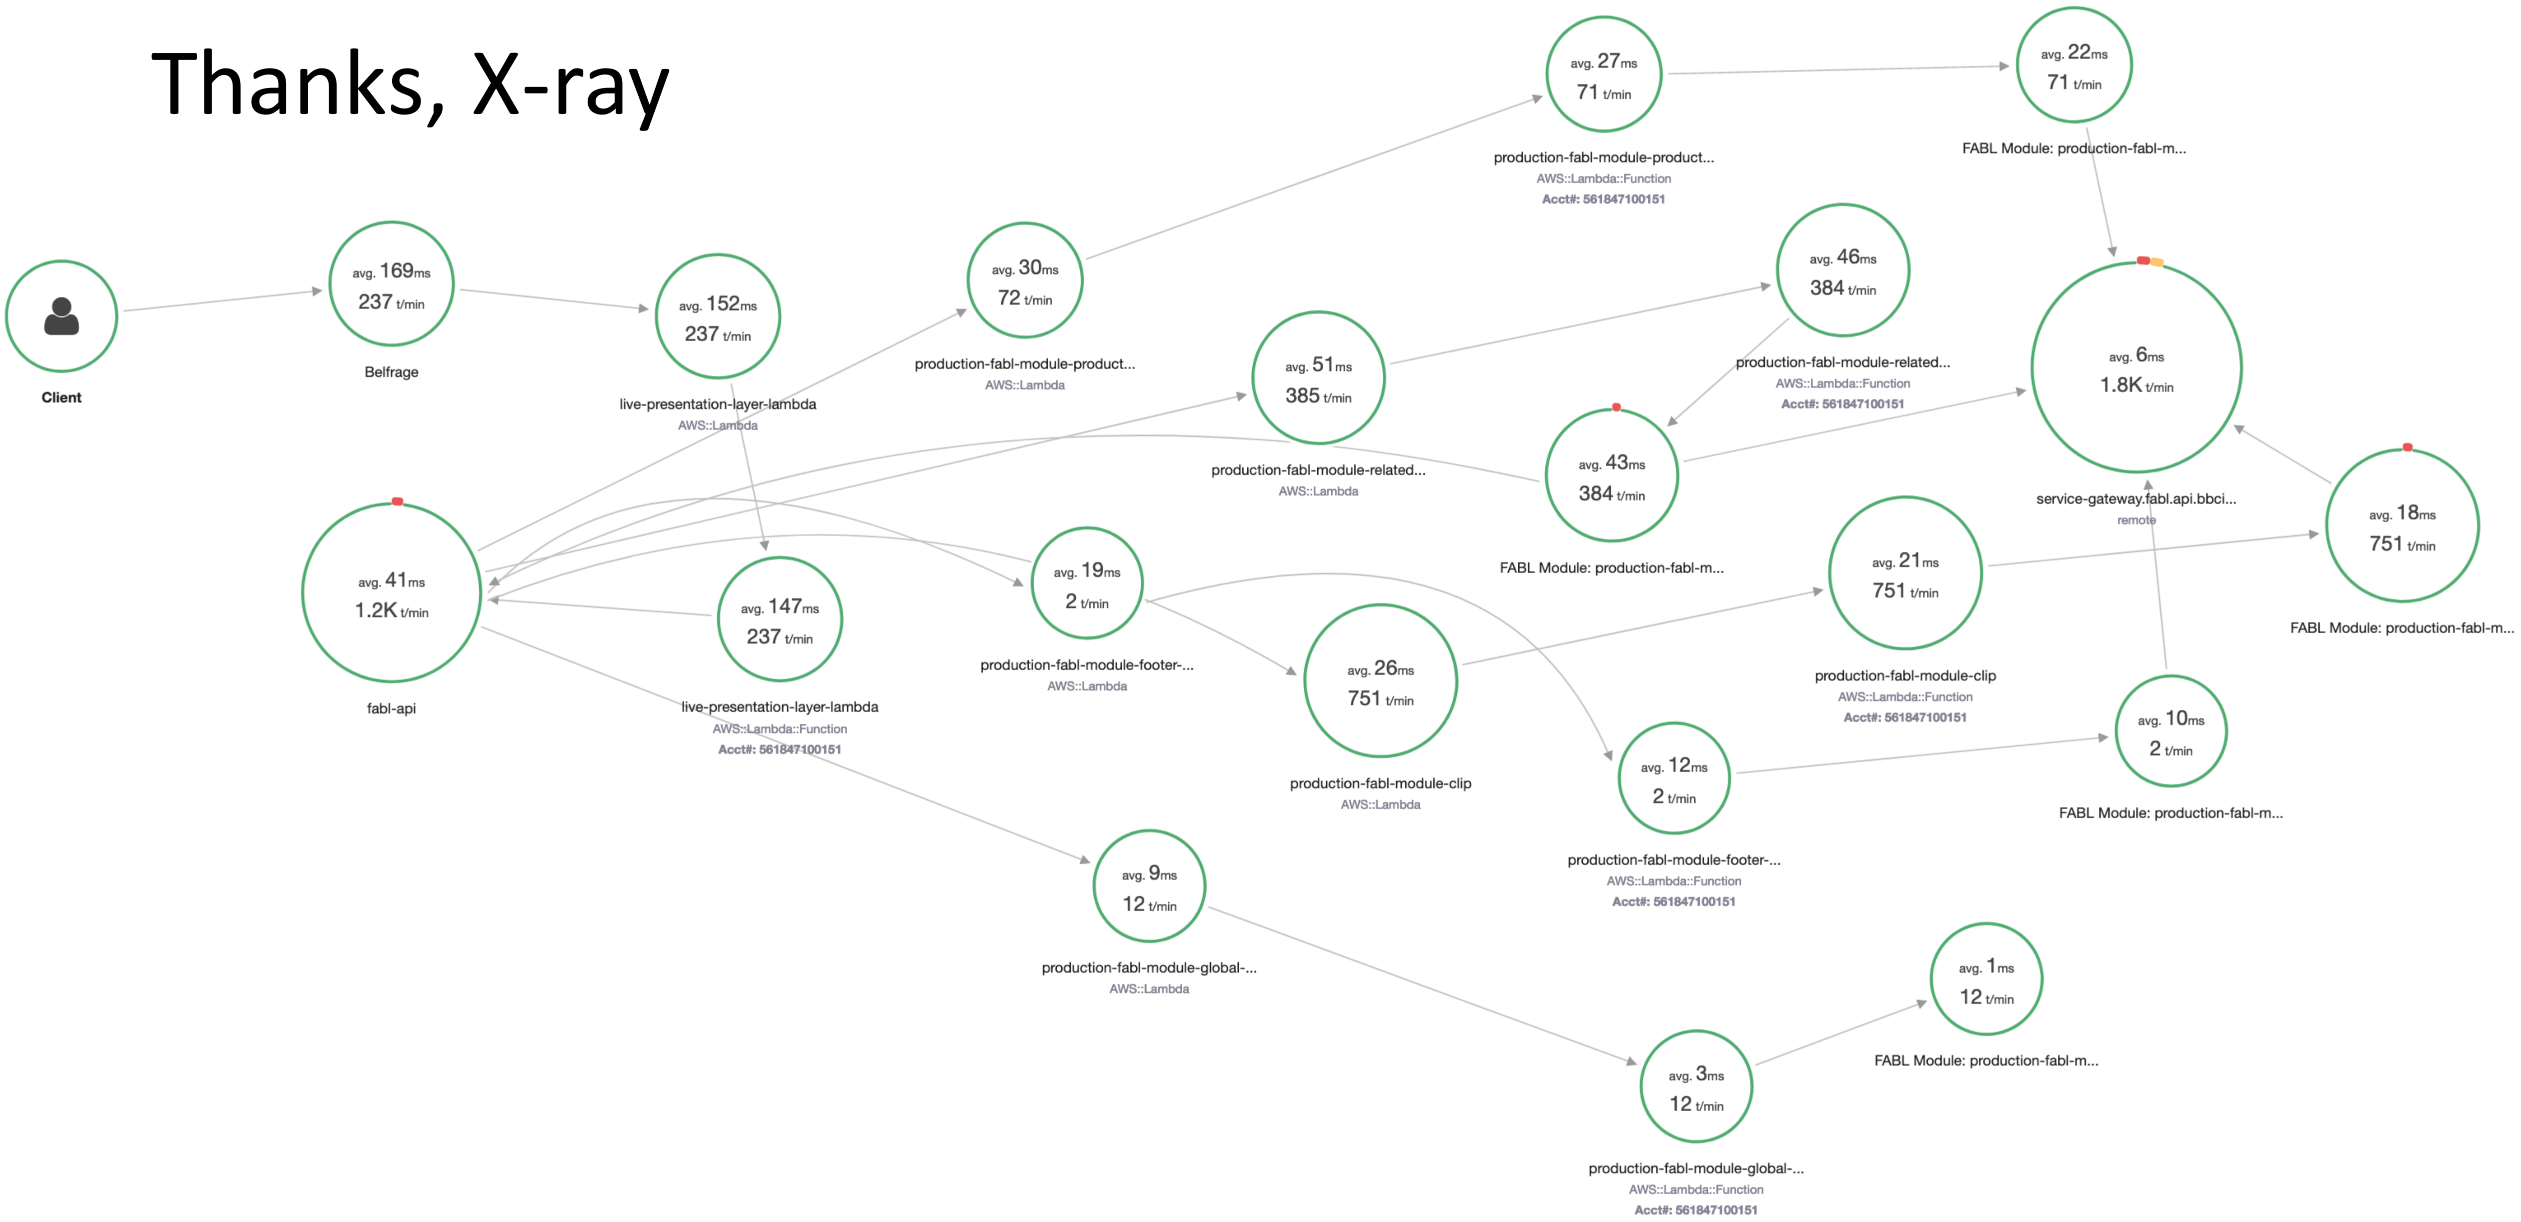
Client
avg. 169ms
237 t/min
Belfrage
avg. 152ms
237 t/min
live-presentation-layer-lambda
AWS::Lambda
avg. 30ms
72 t/min
production-fabl-module-product...
AWS::Lambda
avg. 27ms
71 t/min
production-fabl-module-product...
AWS::Lambda::Function
Acct#: 561847100151
avg. 22ms
71 t/min
FABL Module: production-fabl-m...
avg. 51ms
385 t/min
production-fabl-module-related...
AWS::Lambda
avg. 46ms
384 t/min
production-fabl-module-related...
AWS::Lambda::Function
Acct#: 561847100151
avg. 43ms
384 t/min
FABL Module: production-fabl-m...
avg. 6ms
1.8K t/min
service-gateway.fabl.api.bbci...
remote
avg. 18ms
751 t/min
FABL Module: production-fabl-m...
avg. 41ms
1.2K t/min
fabl-api
avg. 147ms
237 t/min
live-presentation-layer-lambda
AWS::Lambda::Function
Acct#: 561847100151
avg. 19ms
2 t/min
production-fabl-module-footer-...
AWS::Lambda
avg. 26ms
751 t/min
production-fabl-module-clip
AWS::Lambda
avg. 21ms
751 t/min
production-fabl-module-clip
AWS::Lambda::Function
Acct#: 561847100151
avg. 12ms
2 t/min
production-fabl-module-footer-...
AWS::Lambda::Function
Acct#: 561847100151
avg. 10ms
2 t/min
FABL Module: production-fabl-m...
avg. 9ms
12 t/min
production-fabl-module-global-...
AWS::Lambda
avg. 3ms
12 t/min
production-fabl-module-global-...
AWS::Lambda::Function
Acct#: 561847100151
avg. 1ms
12 t/min
FABL Module: production-fabl-m...

Multiple teams → inconsistency

```
"title": "Pengiuns in Antarctic set up marching band"
"headline": "Pengiuns in Antarctic set up marching band"
```

```
"timestamp": 1648399482000
"firstPublished": "2022-03-27T11:49:00.000Z"
```

```
"link": "https://www.bbc.co.uk/news/uk-wales-60589511"
"link":{
    "url": "https://www.bbc.co.uk/news/uk-wales-60589511"
}
```

→ Have conventions

Designing APIs for **change**

IBL – Product as an API

The iPlayer Business Layer (IBL) API is...

- The source of truth for all Product concepts
- The source of truth for content object presence and order
- Not opinionated about the UI

The personalised version of the API is exposed over GraphQL

IBL – Product as an API

Every property exposed has a cost once published

Image - https://www.flickr.com/photos/blmnevada/15519491237/in/photolist-pDps48

IBL – Product as an API

- Know the purpose of your API
- Design explicitly for your clients use cases
- Limit your supported consumers to those use cases
- Persisted GraphQL queries can help manage evolution!

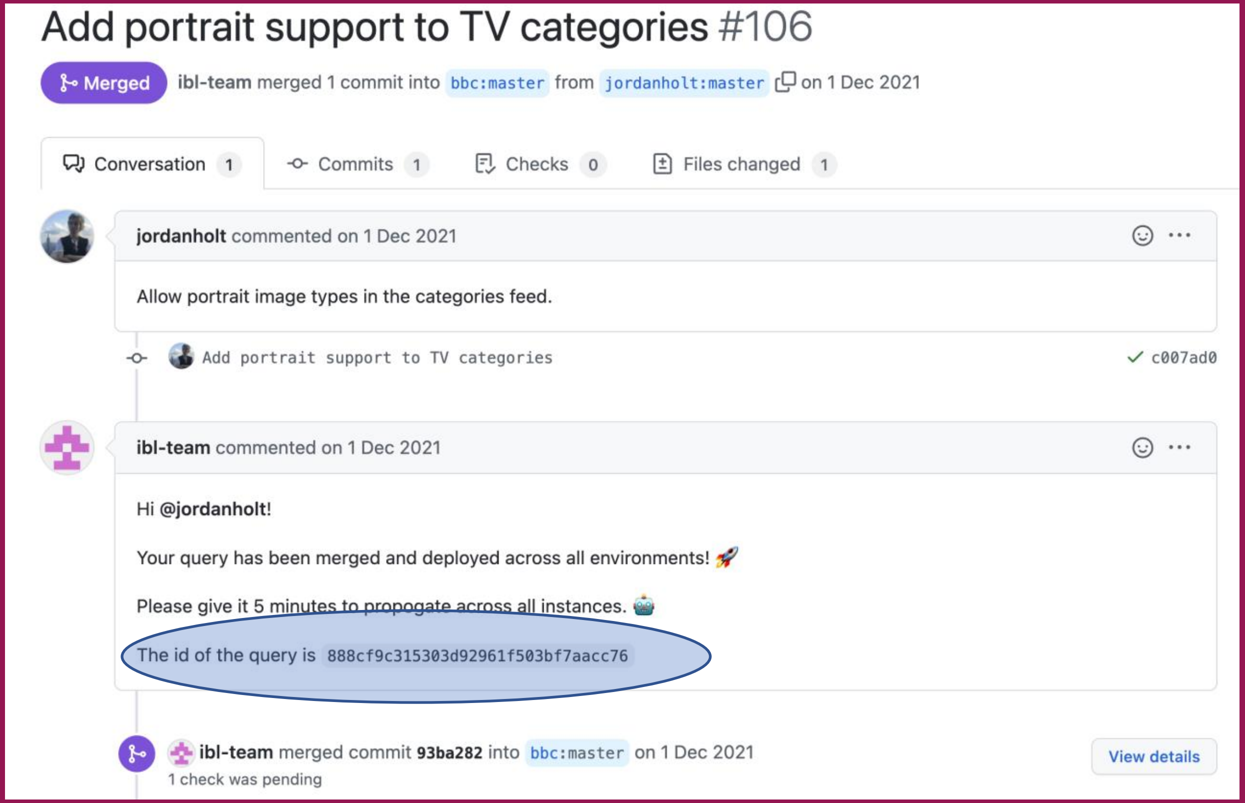

Add portrait support to TV categories #106

ibl-team merged 1 commit into `bbc:master` from `jordanholt:master` on 1 Dec 2021

jordanholt commented on 1 Dec 2021

Allow portrait image types in the categories feed.

```
Add portrait support to TV categories
```
✓ c007ad0

ibl-team commented on 1 Dec 2021

Hi @jordanholt!

Your query has been merged and deployed across all environments! 🚀

Please give it 5 minutes to propogate across all instances. 🤖

The id of the query is `888cf9c315303d92961f503bf7aacc76`

ibl-team merged commit 93ba282 into `bbc:master` on 1 Dec 2021

1 check was pending

Handling API change

```
{
    "id": "urn:bbc:article:123",
    "url": "https://www.bbc.com/news/articles/123",
    "firstPublished": "2022-03-27T11:49:00.000Z",
    "title": "Duck swims around the world",
    "category": "nature",
    "categories": ["nature", "travel"],
}
```

```
<image deprecated="true" deprecated_since="2014-07-09"
replaced_by="mixin=images" url="http://..."/>
```

Actual BBC API example

```
image: String @deprecated(reason:
"image is deprecated. Use imageUrl instead.")
```

GraphQL schema

API versioning

Version the whole API…

```
https://api.twitter.com/1.1/statuses/lookup.json
https://api.twitter.com/2/tweets
```

Or just one end-point…

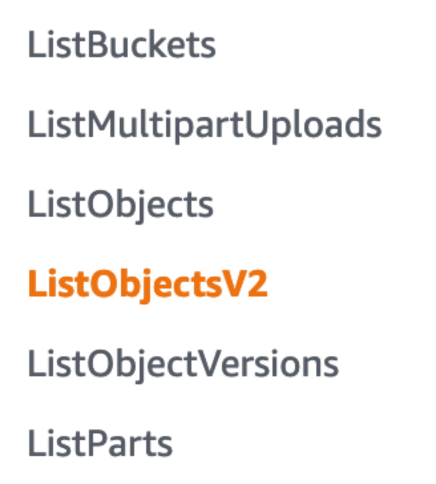

ListBuckets

ListMultipartUploads

ListObjects

ListObjectsV2

ListObjectVersions

ListParts

Example from AWS S3 API

Many clients

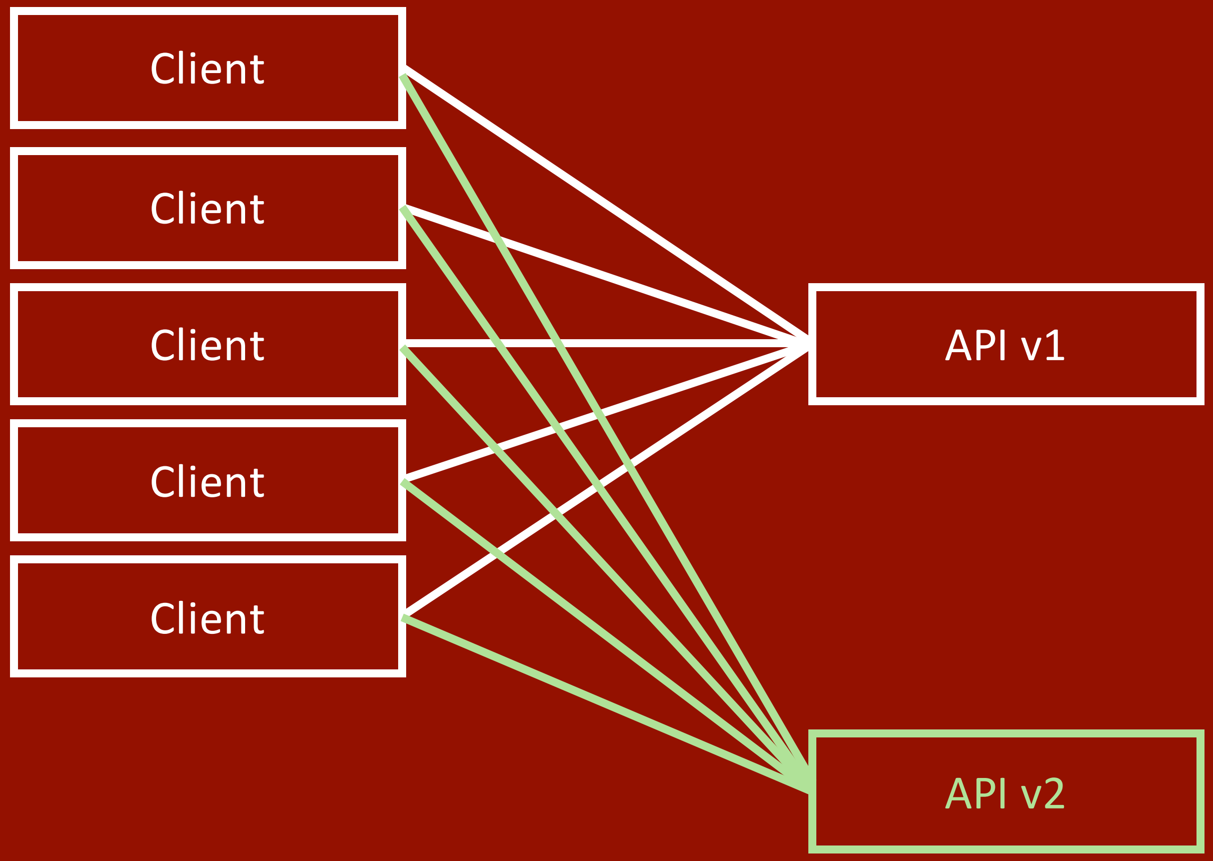

From: Joe Developer

To: User d-list

Subject: New API Version

You must upgrade!

The old one will be switched off, July 1

Don't worry!

The original API will continue working indefinitely.

Creating a facade

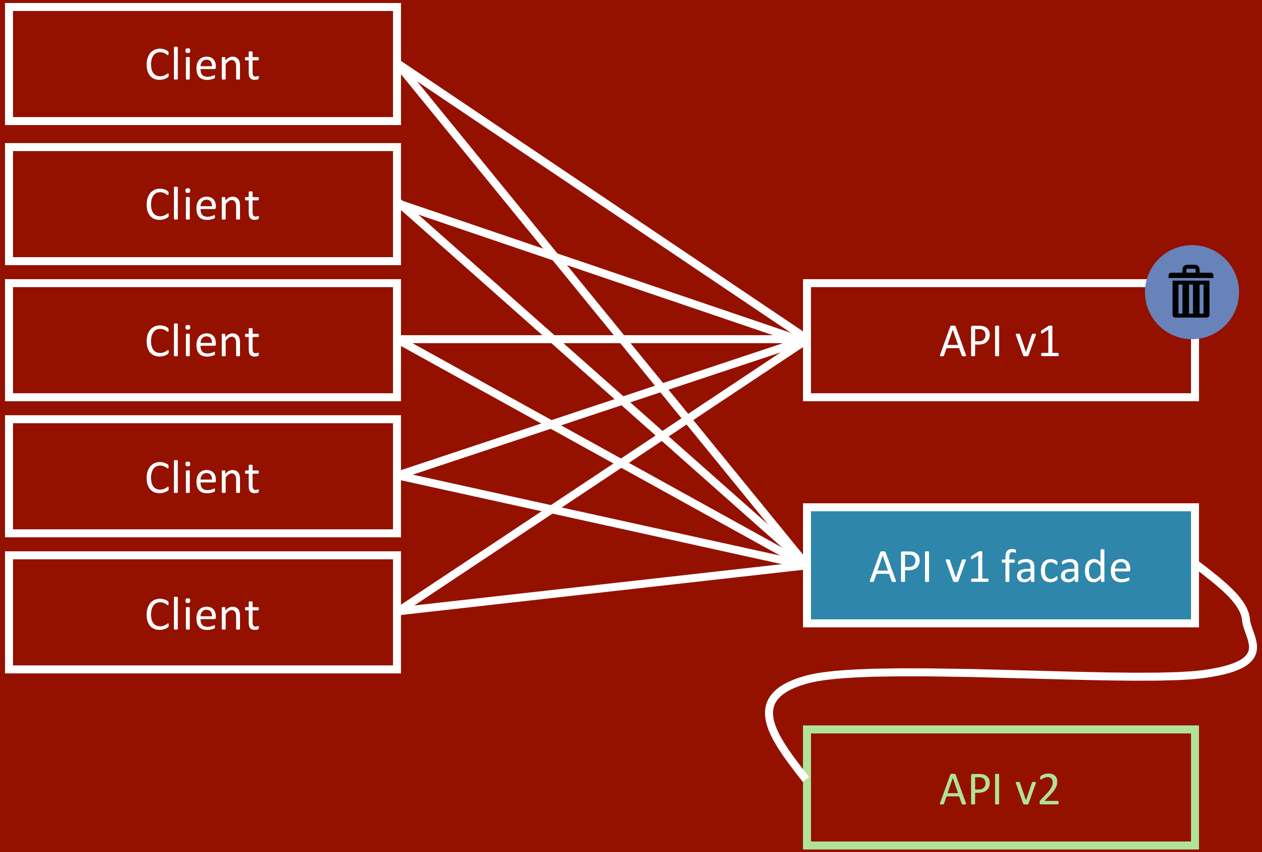

Know who’s involved

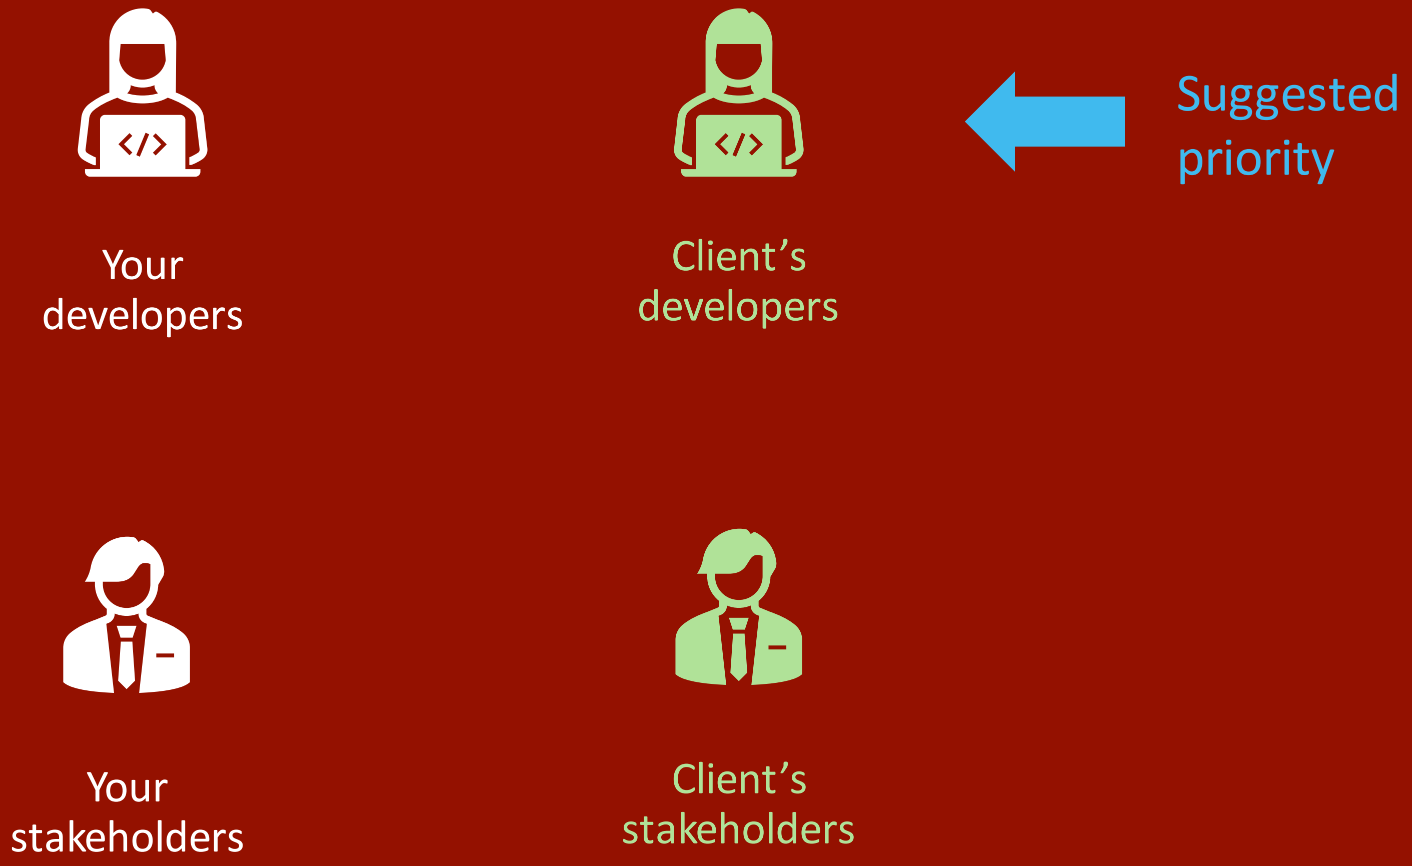

Your developers

Client’s developers

Suggested priority

Your stakeholders

Client’s stakeholders

Designing APIs for **scale**

A broadcast Denial of Service attack

Serverless Launch

- Leverage CDN edge functions
- Statically pre-published responses per device to S3
- Device categorization in lambda@edge
- Augment with dynamic values
- But what about the next hop…?

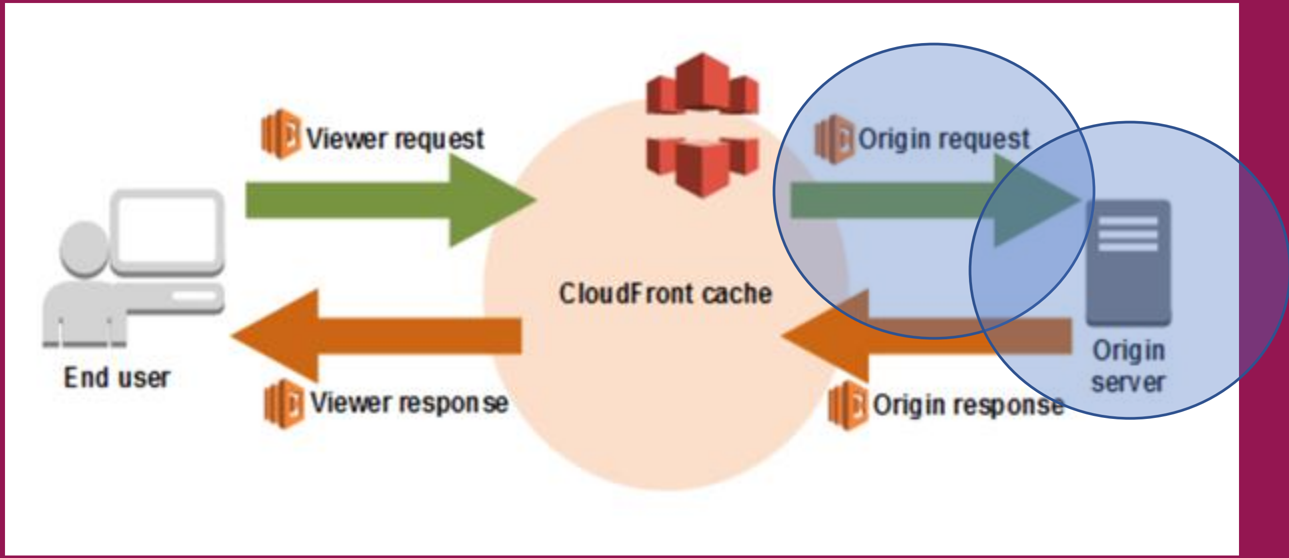

Source: https://docs.aws.amazon.com/lambda/latest/dg/lambda-edge.html

Fallbacks - failure as a first class citizen

- Fallbacks - planning for failure as a first class part of the API design, not an afterthought
- What do fallbacks need to do?
  - Provides schema compliant lighter responses
  - Supports context sensitivity
  - Enable scale by offloading
- Trade offs, trade offs everywhere
  - Availability over degraded features and timeliness of content
  - Polling vs event driven publishing

Serving Fallbacks

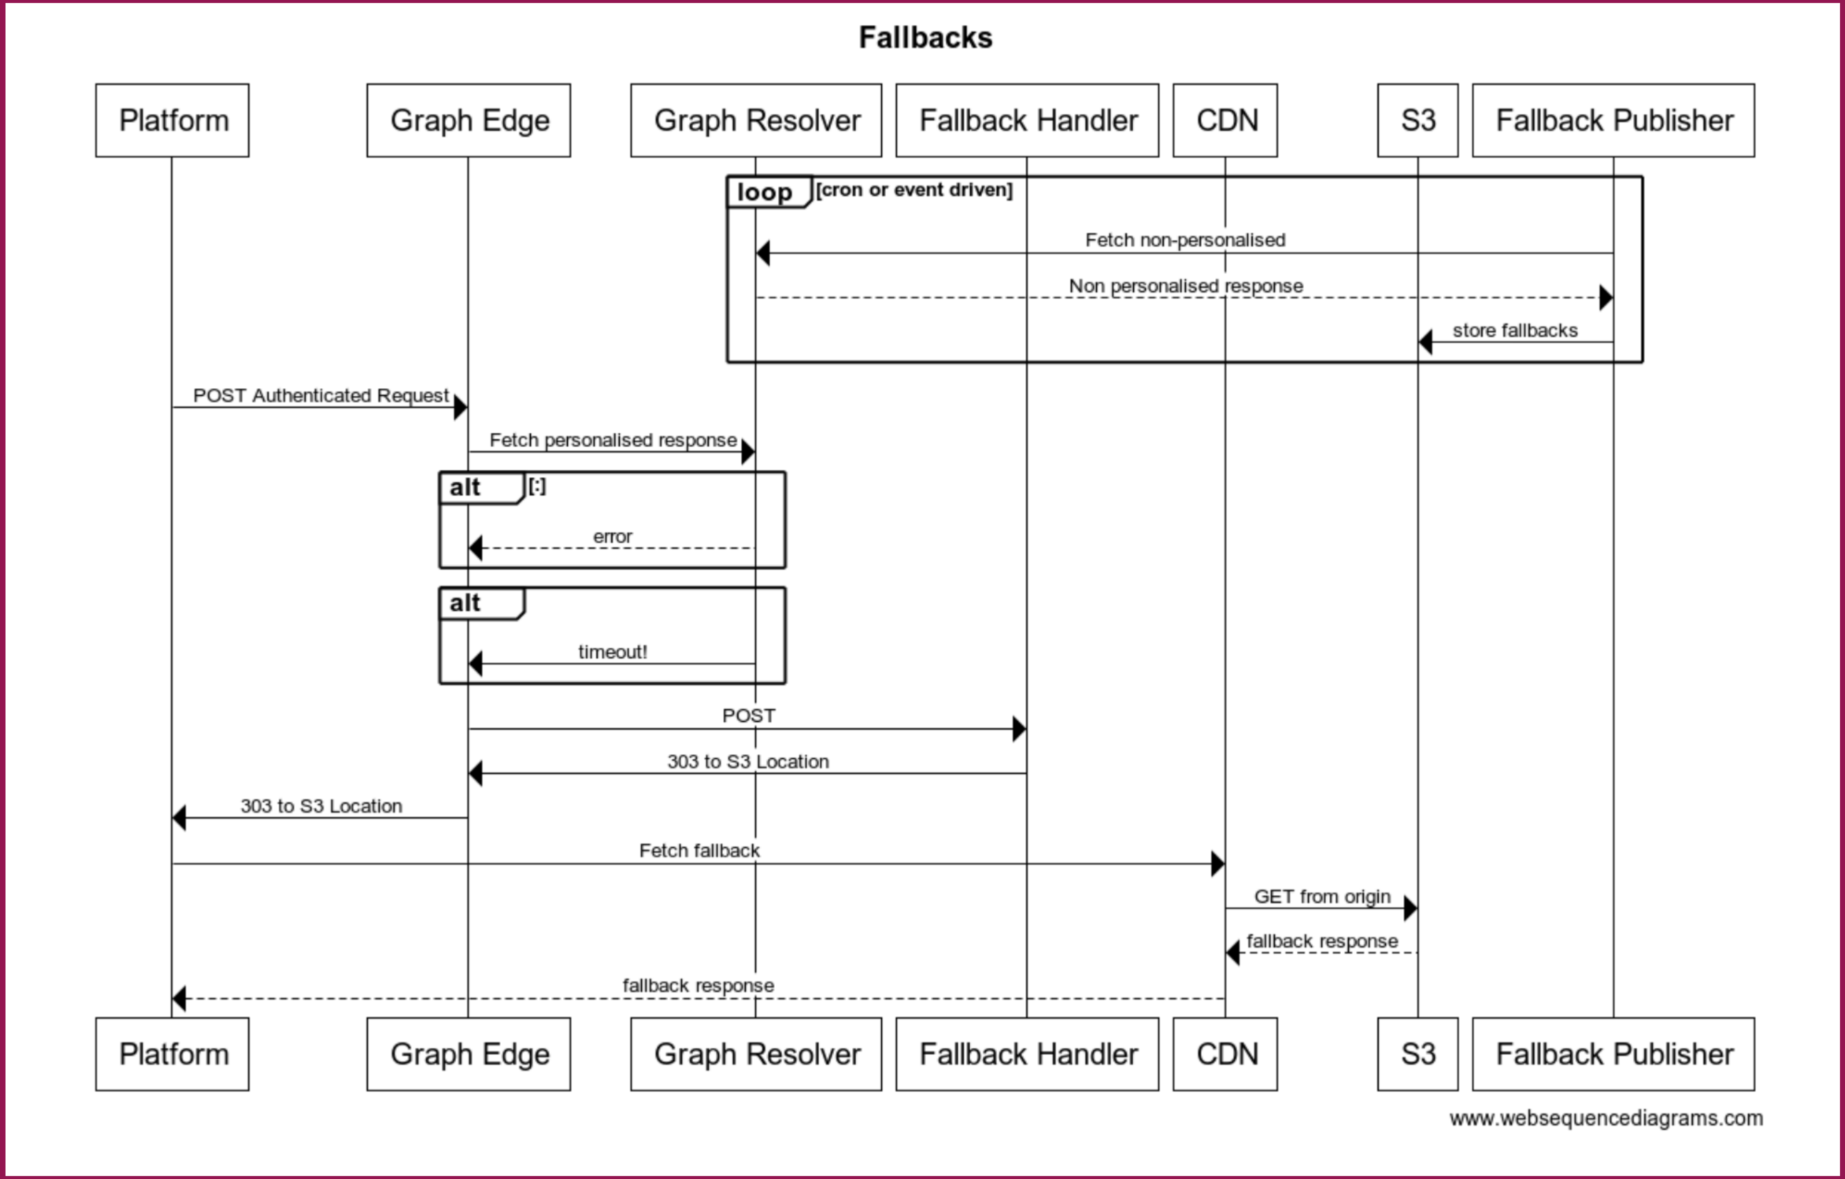

Fallbacks

Platform

Graph Edge

Graph Resolver

Fallback Handler

CDN

S3

Fallback Publisher

loop [cron or event driven]

Fetch non-personalised

Non personalised response

store fallbacks

POST Authenticated Request

Fetch personalised response

alt [:]

error

alt

timeout!

POST

303 to S3 Location

303 to S3 Location

Fetch fallback

GET from origin

fallback response

fallback response

Platform

Graph Edge

Graph Resolver

Fallback Handler

CDN

S3

Fallback Publisher

www.websequencediagrams.com

Serving Fallbacks

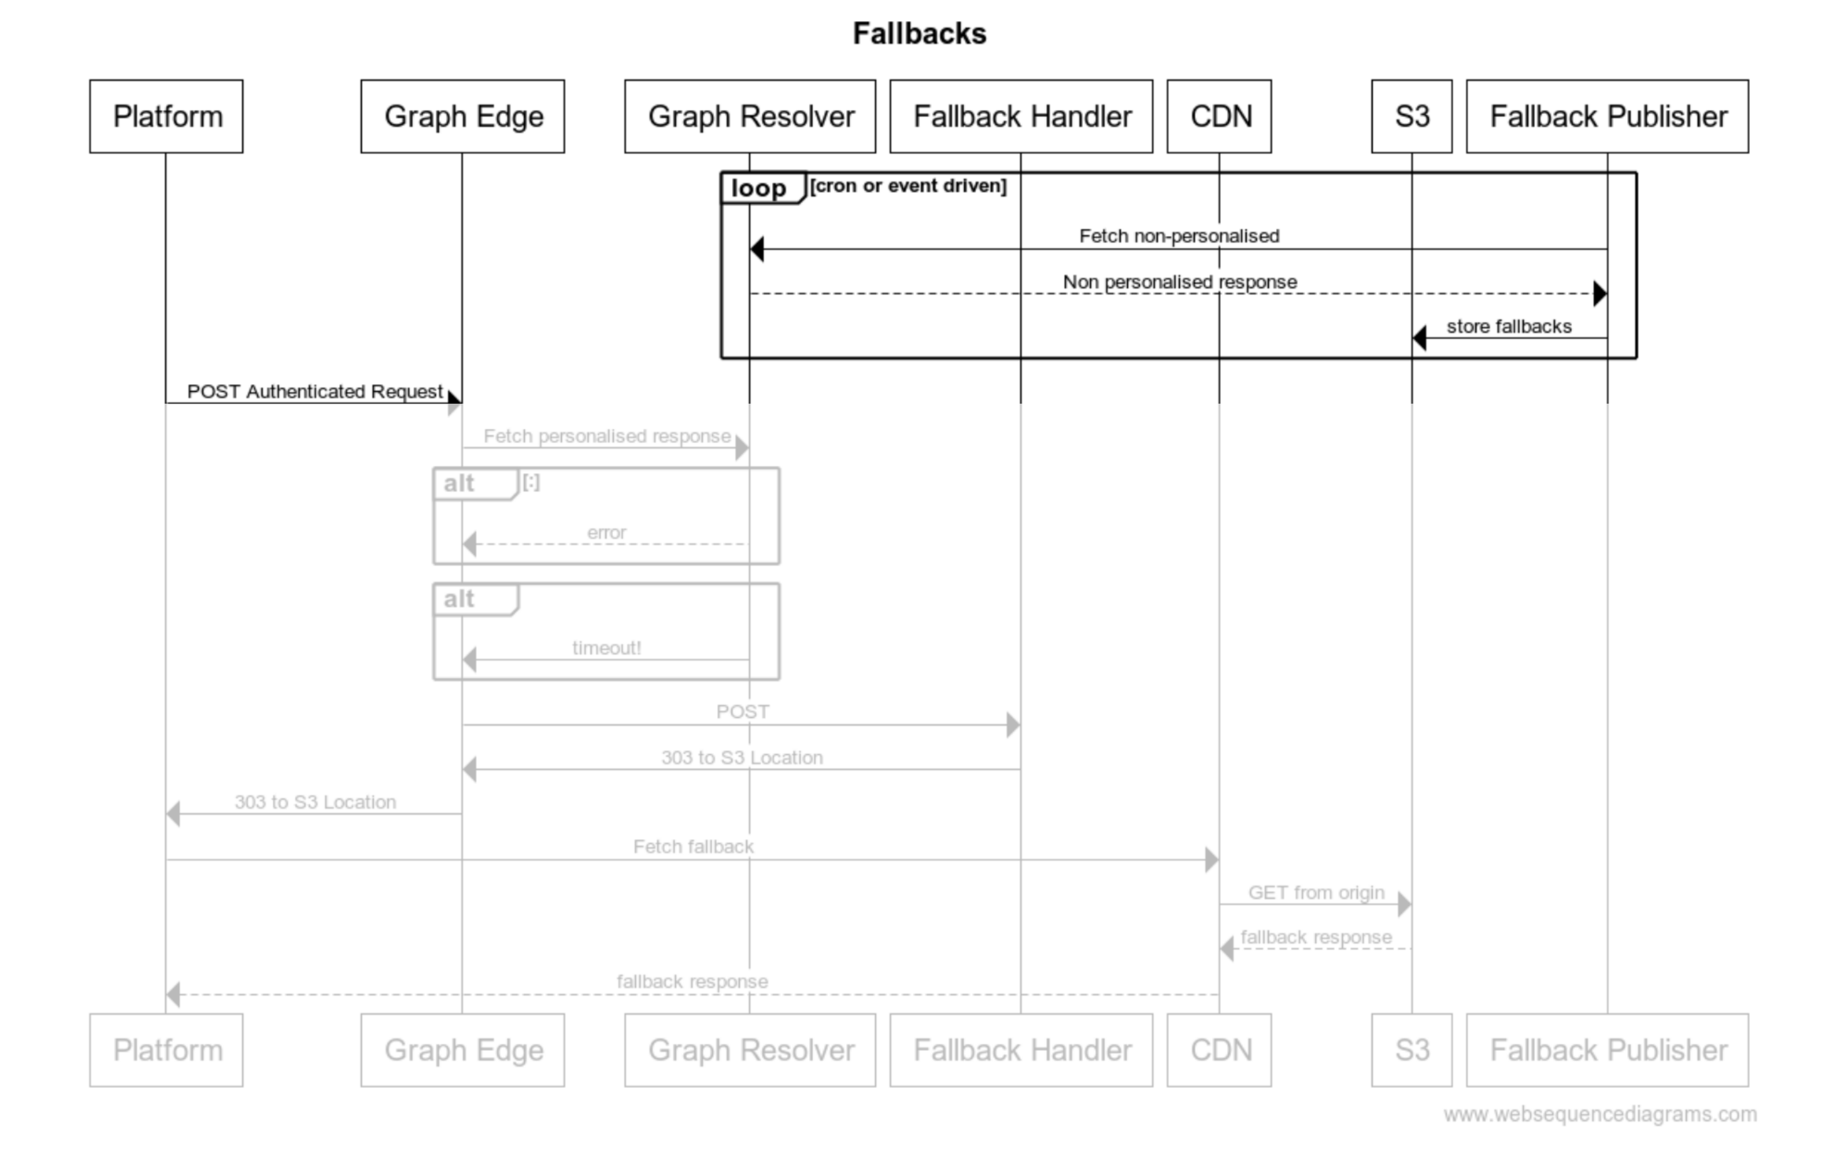

Fallbacks

Platform

Graph Edge

Graph Resolver

Fallback Handler

CDN

S3

Fallback Publisher

loop [cron or event driven]

Fetch non-personalised

Non personalised response

store fallbacks

POST Authenticated Request

Fetch personalised response

alt [:]

error

alt

timeout!

POST

303 to S3 Location

303 to S3 Location

Fetch fallback

GET from origin

fallback response

fallback response

Platform

Graph Edge

Graph Resolver

Fallback Handler

CDN

S3

Fallback Publisher

www.websequencediagrams.com

Serving Fallbacks

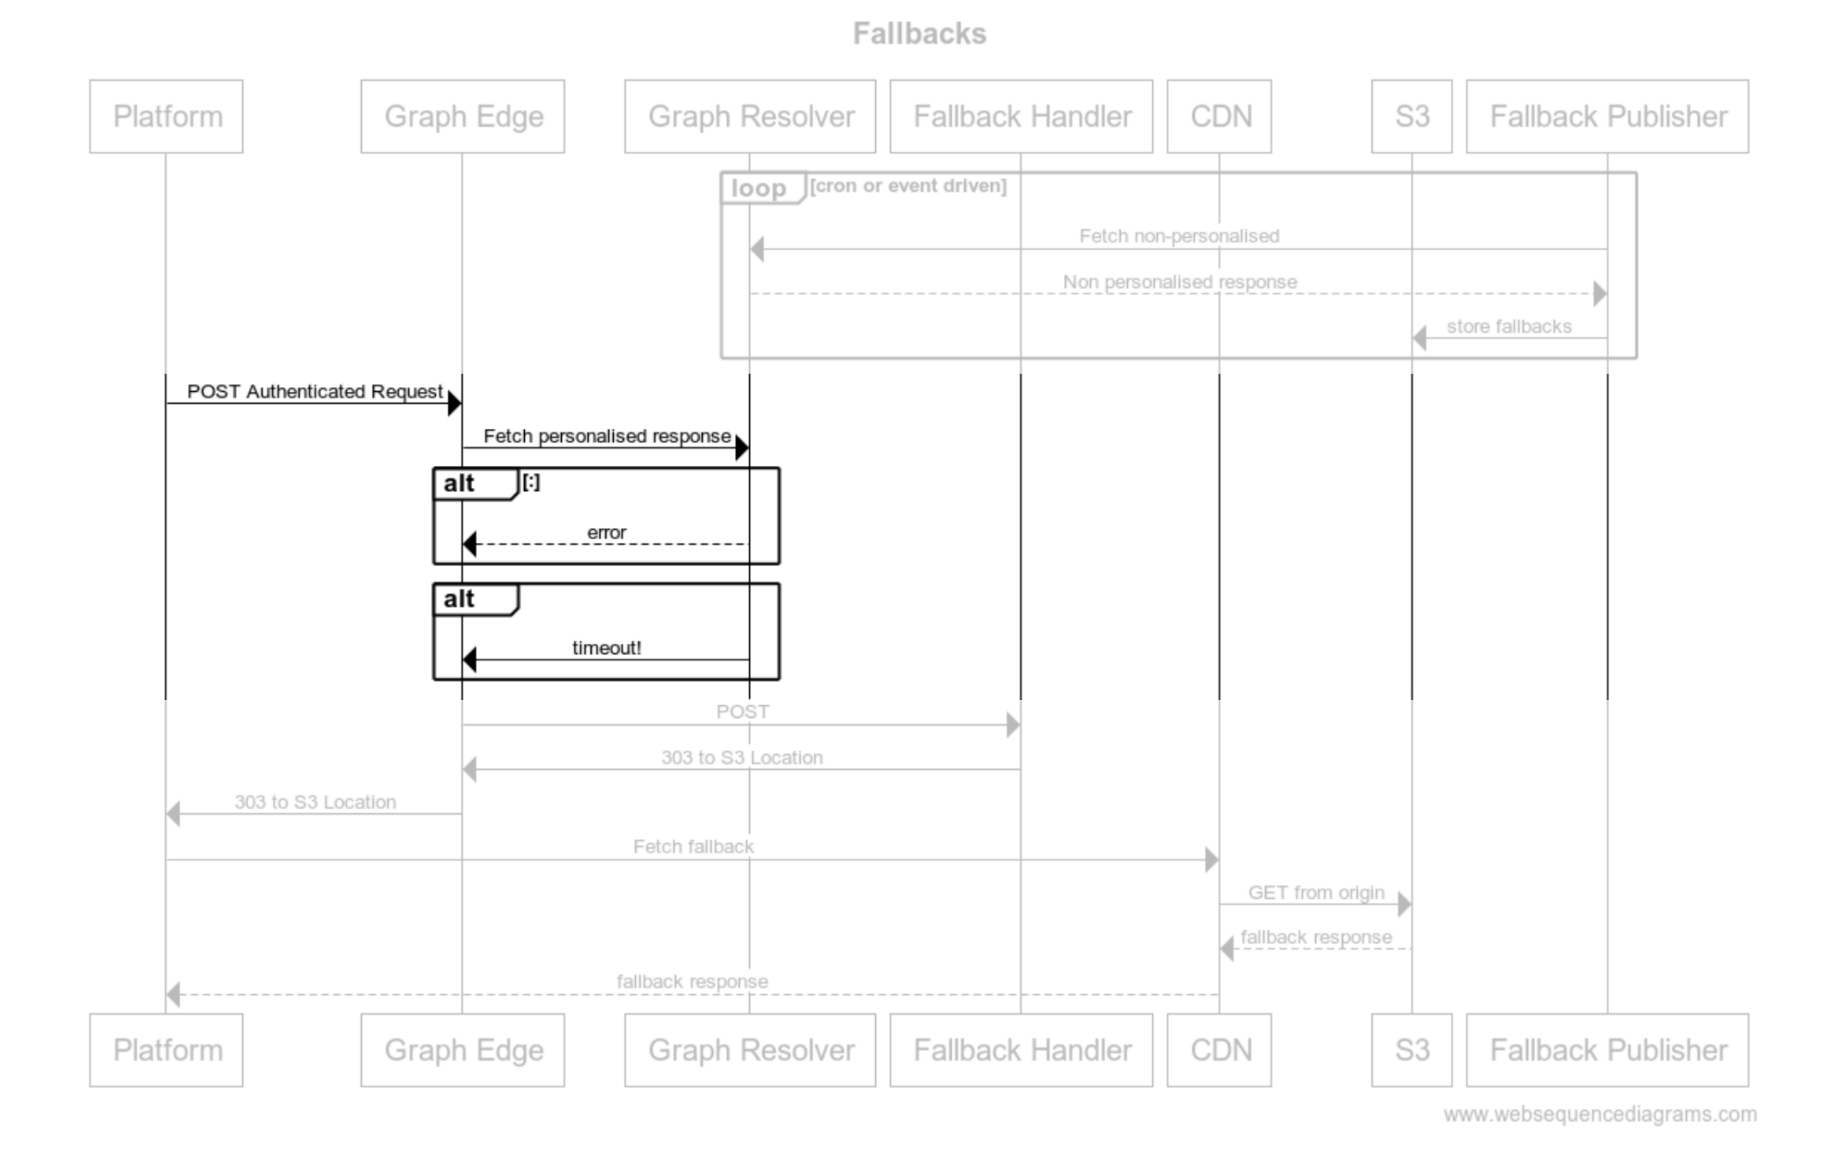

Serving Fallbacks

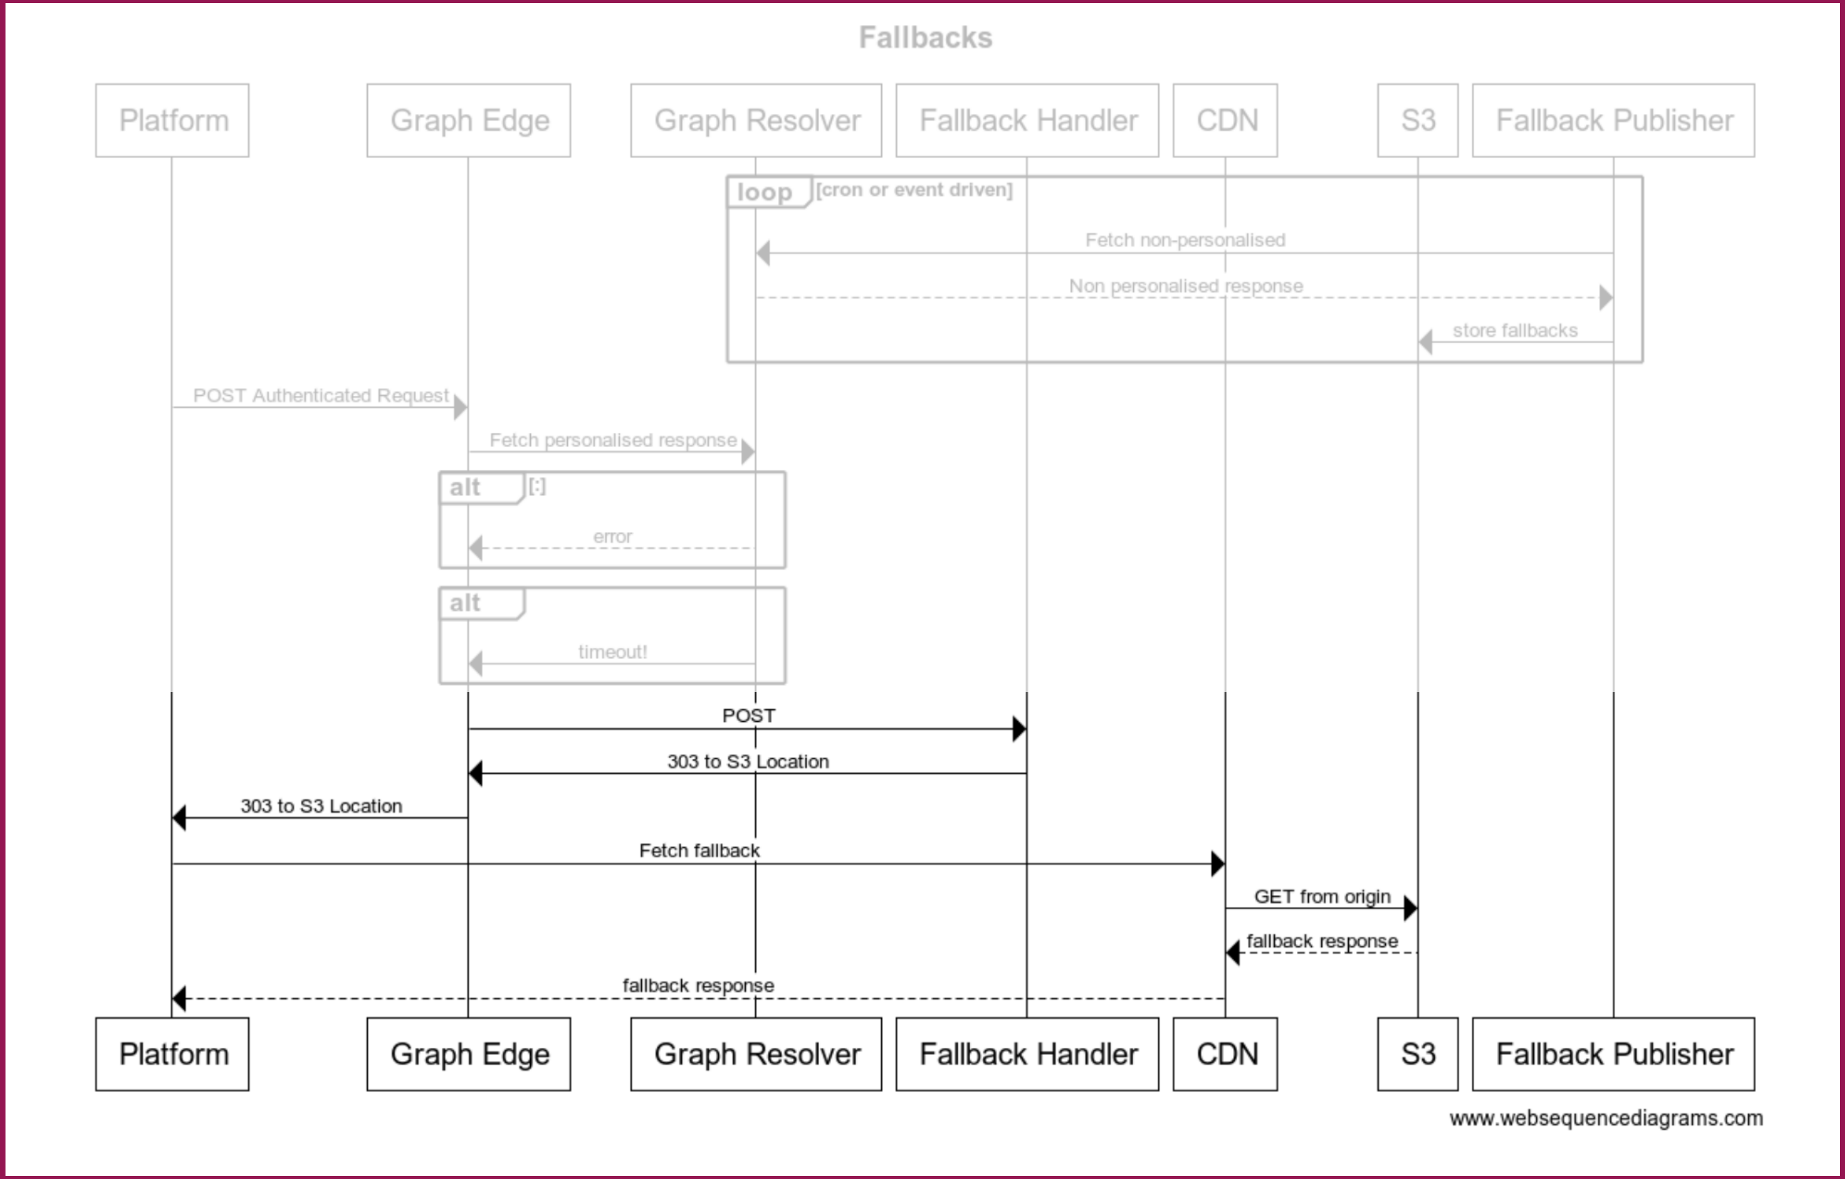

Serving Fallbacks

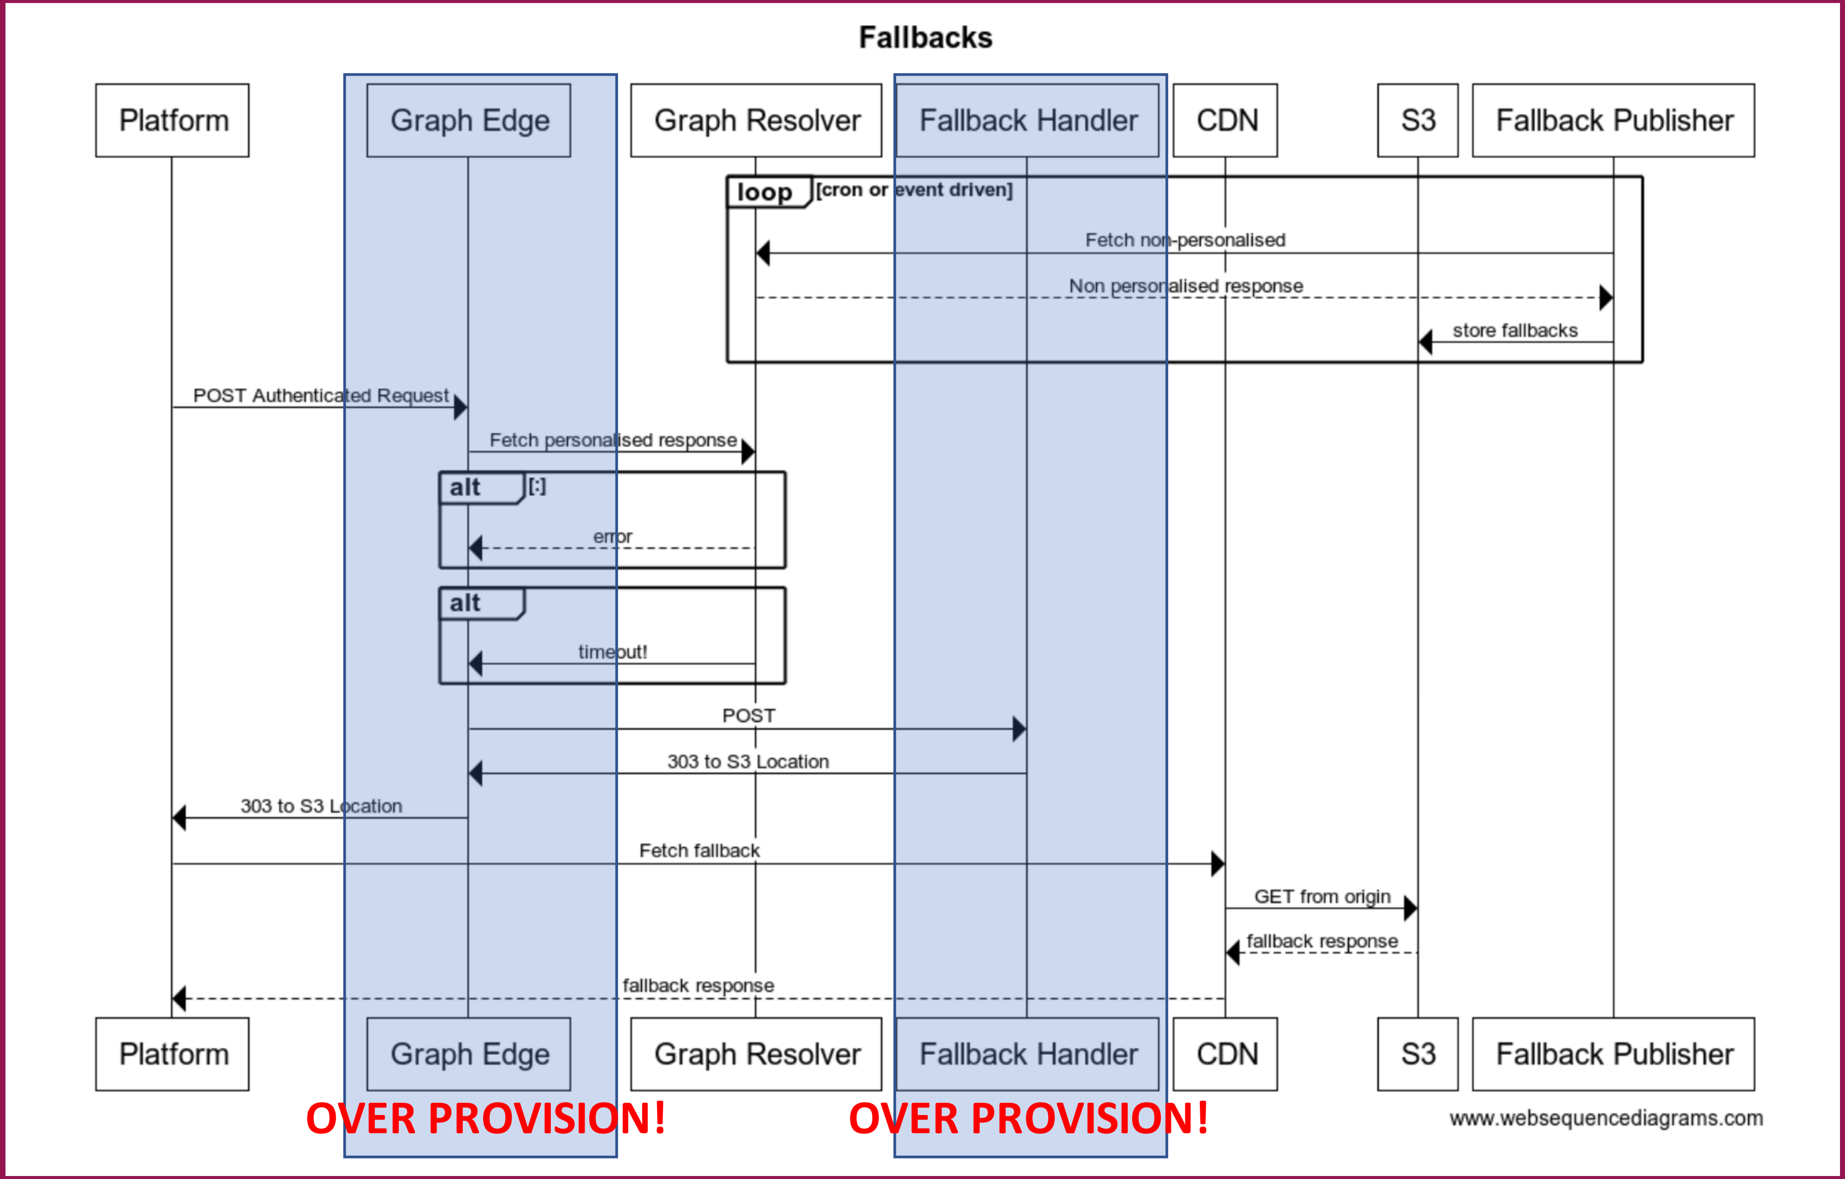

A third of the BBC's pages are rendered using serverless (including the APIs that power them)

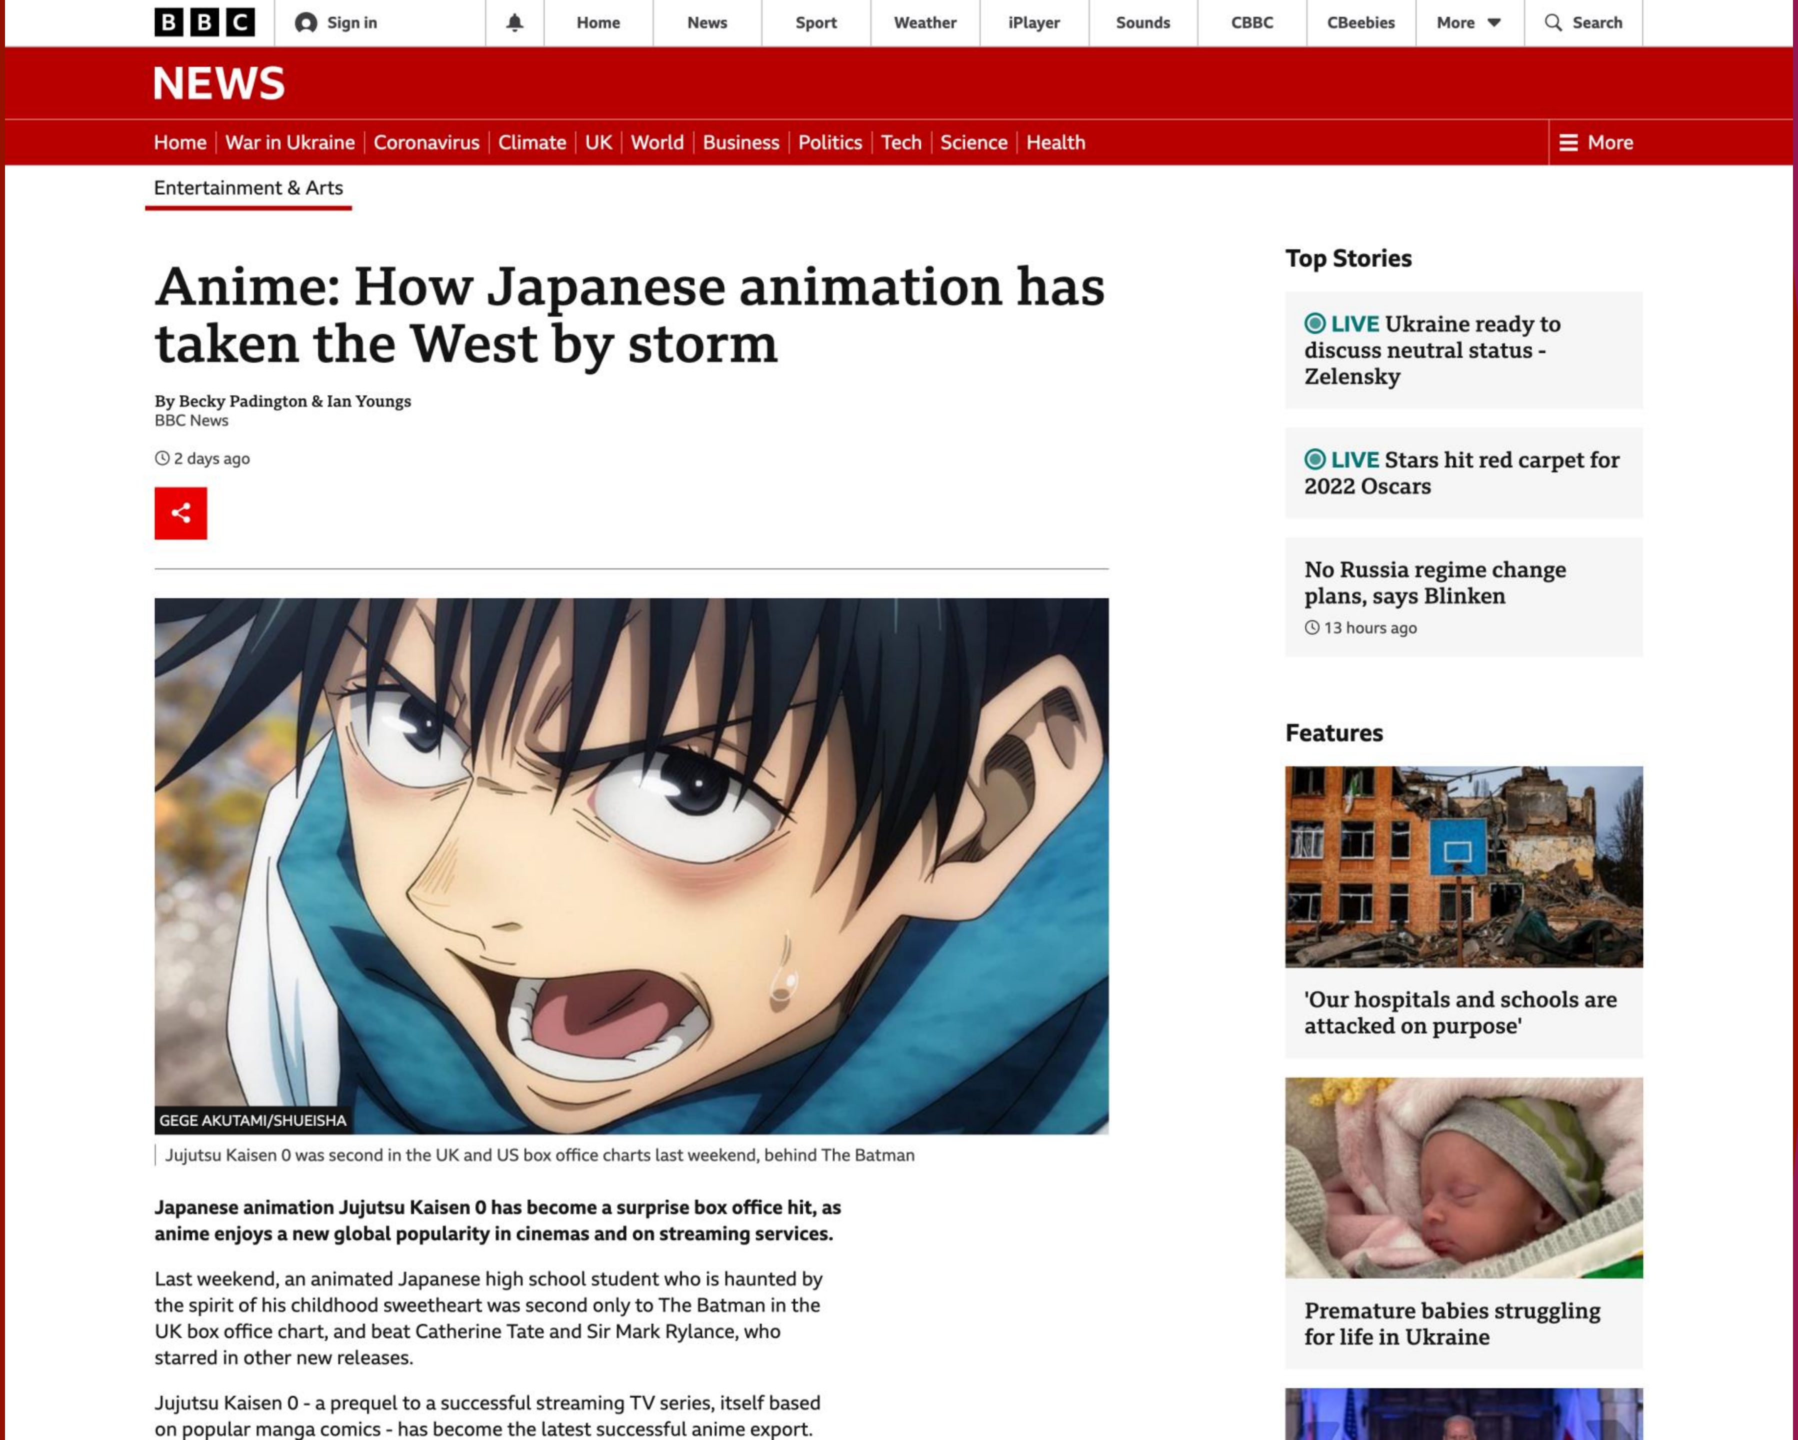

BBC
Sign in
Home
News
Sport
Weather
iPlayer
Sounds
CBBC
CBeebies
More
Search
NEWS
Home | War in Ukraine | Coronavirus | Climate | UK | World | Business | Politics | Tech | Science | Health
More
Entertainment & Arts

Anime: How Japanese animation has taken the West by storm

By Becky Padington & Ian Youngs
BBC News

2 days ago

GEGE AKUTAMI/SHUEISHA

Jujutsu Kaisen 0 was second in the UK and US box office charts last weekend, behind The Batman

Japanese animation Jujutsu Kaisen 0 has become a surprise box office hit, as anime enjoys a new global popularity in cinemas and on streaming services.

Last weekend, an animated Japanese high school student who is haunted by the spirit of his childhood sweetheart was second only to The Batman in the UK box office chart, and beat Catherine Tate and Sir Mark Rylance, who starred in other new releases.

Jujutsu Kaisen 0 - a prequel to a successful streaming TV series, itself based on popular manga comics - has become the latest successful anime export.

Top Stories

LIVE Ukraine ready to discuss neutral status - Zelensky

LIVE Stars hit red carpet for 2022 Oscars

No Russia regime change plans, says Blinken
13 hours ago

Features

'Our hospitals and schools are attacked on purpose'

Premature babies struggling for life in Ukraine

If using AWS, the CDK patterns make serverless API development straightforward

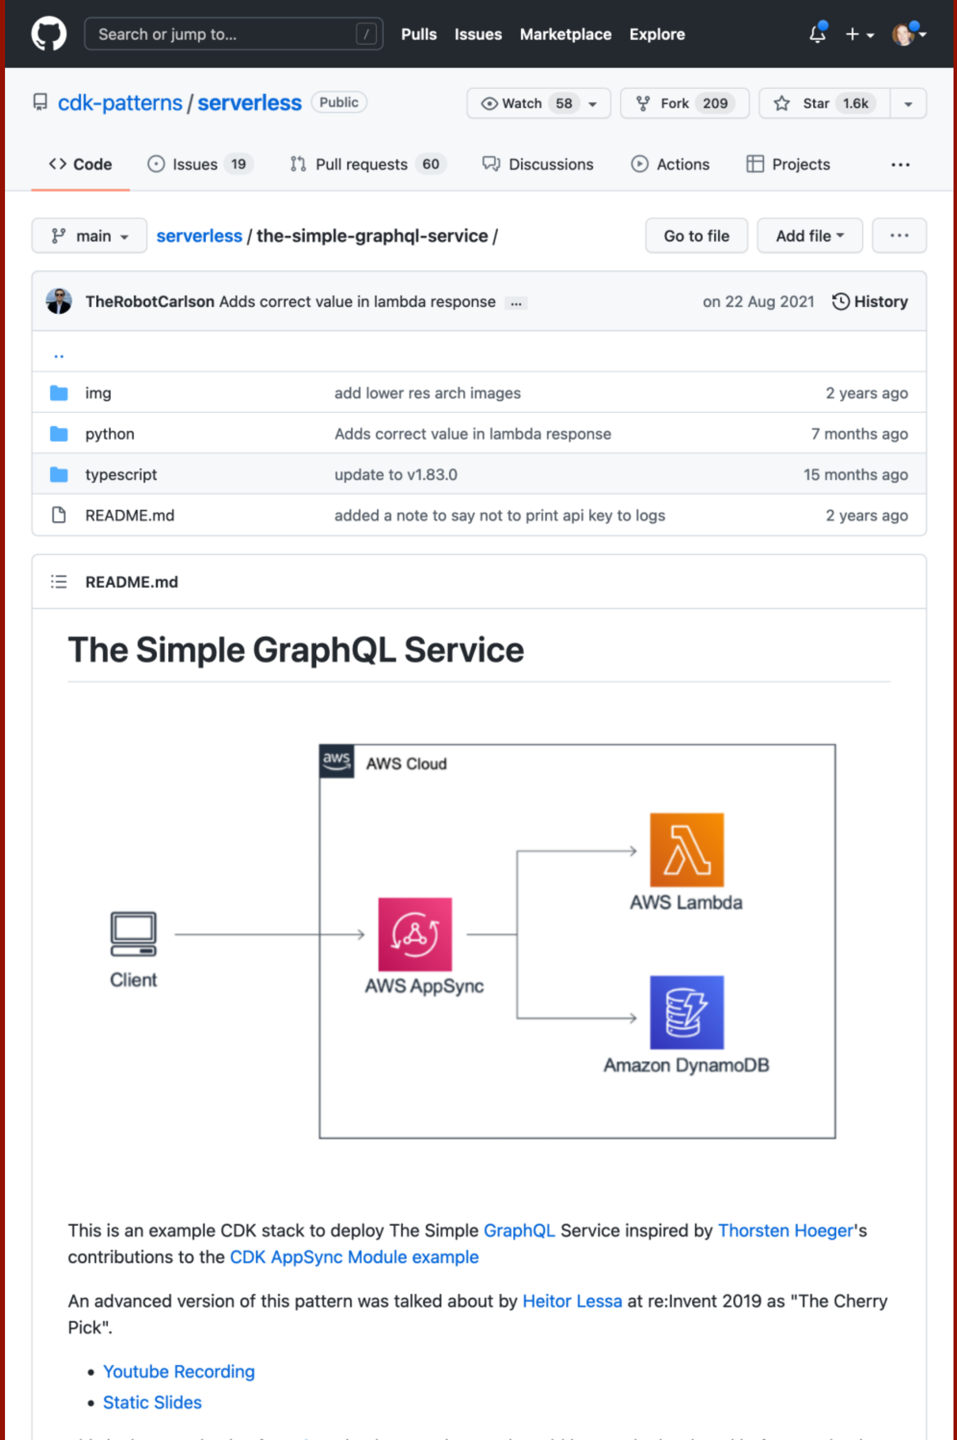

TheRobotCarlson Adds correct value in lambda response ... on 22 Aug 2021

..

img	add lower res arch images	2 years ago
python	Adds correct value in lambda response	7 months ago
typescript	update to v1.83.0	15 months ago
README.md	added a note to say not to print api key to logs	2 years ago

README.md

The Simple GraphQL Service

This is an example CDK stack to deploy The Simple GraphQL Service inspired by Thorsten Hoeger's contributions to the CDK AppSync Module example

An advanced version of this pattern was talked about by Heitor Lessa at re:Invent 2019 as "The Cherry Pick".

- Youtube Recording
- Static Slides

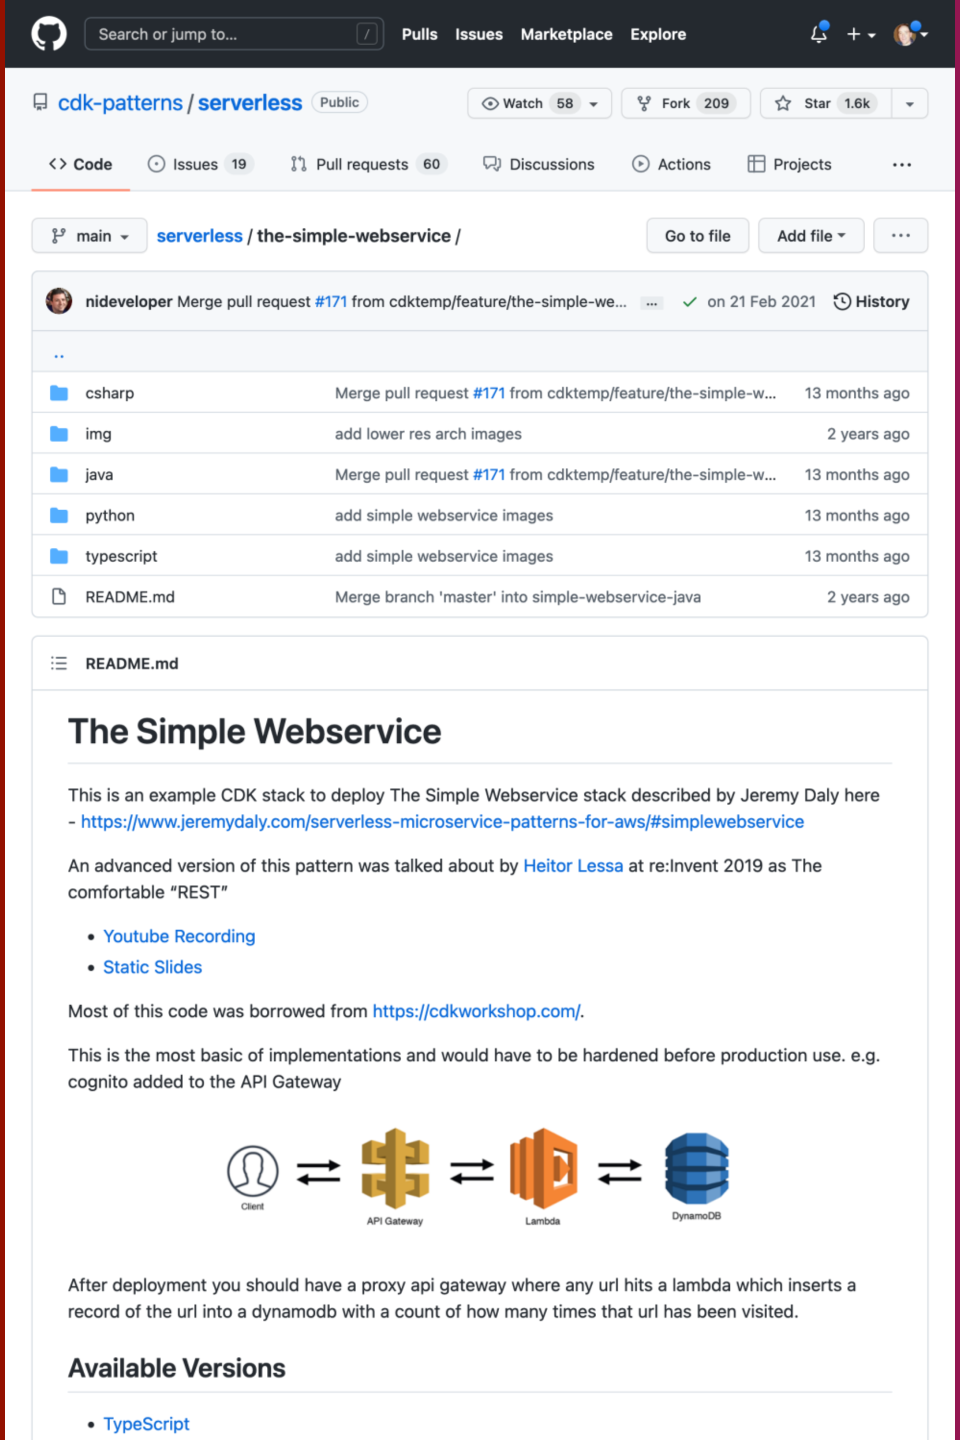

nideveloper Merge pull request #171 from cdktemp/feature/the-simple-we... ... ✓ on 21 Feb 2021

..

csharp	Merge pull request #171 from cdktemp/feature/the-simple-w...	13 months ago
img	add lower res arch images	2 years ago
java	Merge pull request #171 from cdktemp/feature/the-simple-w...	13 months ago
python	add simple webservice images	13 months ago
typescript	add simple webservice images	13 months ago
README.md	Merge branch 'master' into simple-webservice-java	2 years ago

README.md

The Simple Webservice

This is an example CDK stack to deploy The Simple Webservice stack described by Jeremy Daly here - https://www.jeremydaly.com/serverless-microservice-patterns-for-aws/#simplewebservice

An advanced version of this pattern was talked about by Heitor Lessa at re:Invent 2019 as The comfortable "REST"

- Youtube Recording
- Static Slides

Most of this code was borrowed from https://cdkworkshop.com/.

This is the most basic of implementations and would have to be hardened before production use. e.g. cognito added to the API Gateway

After deployment you should have a proxy api gateway where any url hits a lambda which inserts a record of the url into a dynamodb with a count of how many times that url has been visited.

Available Versions

- TypeScript

https://github.com/cdk-patterns/serverless

Low-code / no-code APIs

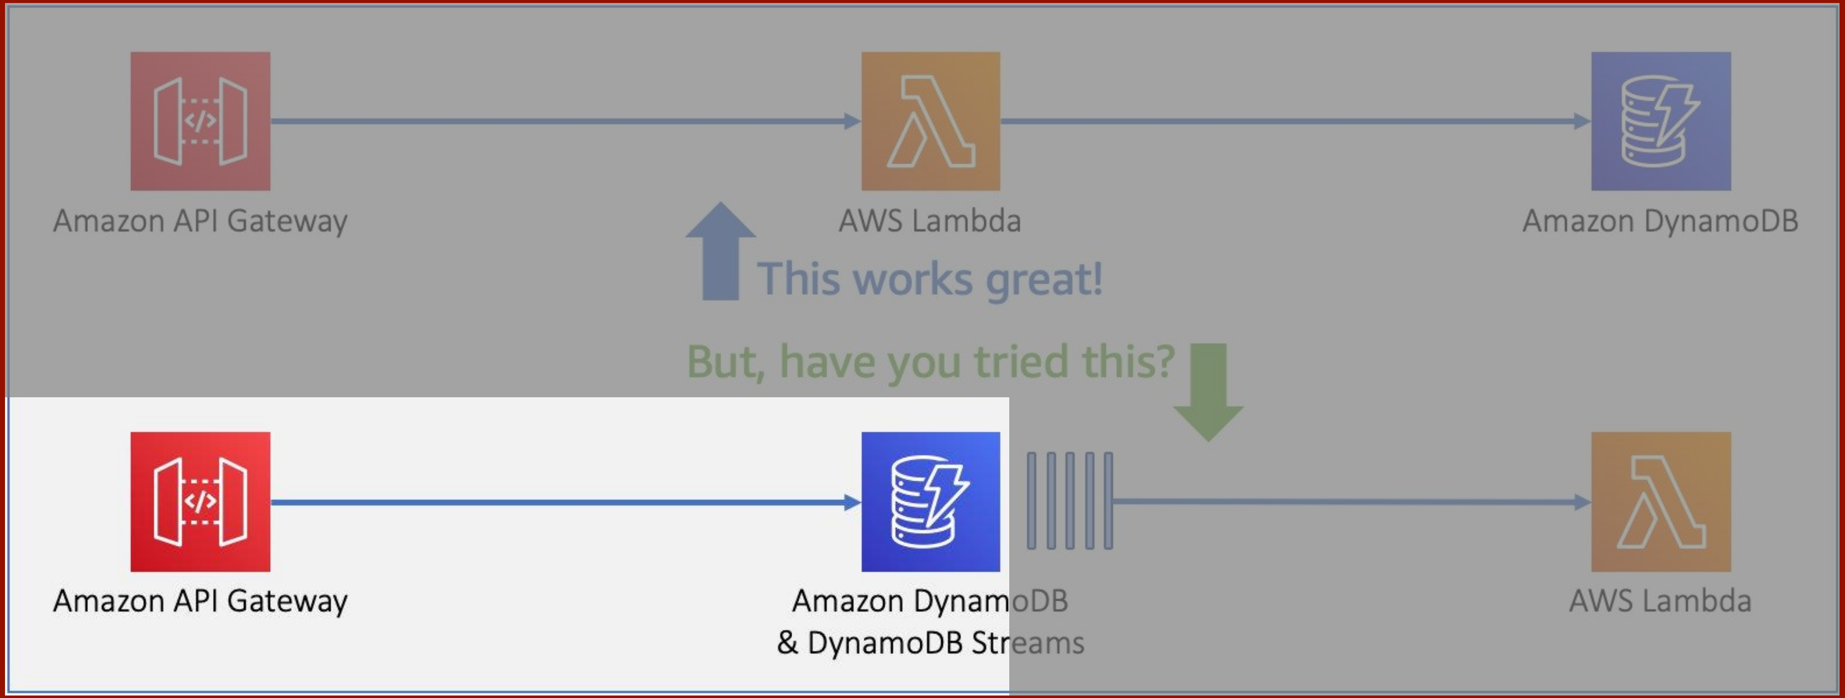

https://github.com/cdk-patterns/serverless/blob/main/the-dynamo-streamer/README.md

Scaling instantly when you need it

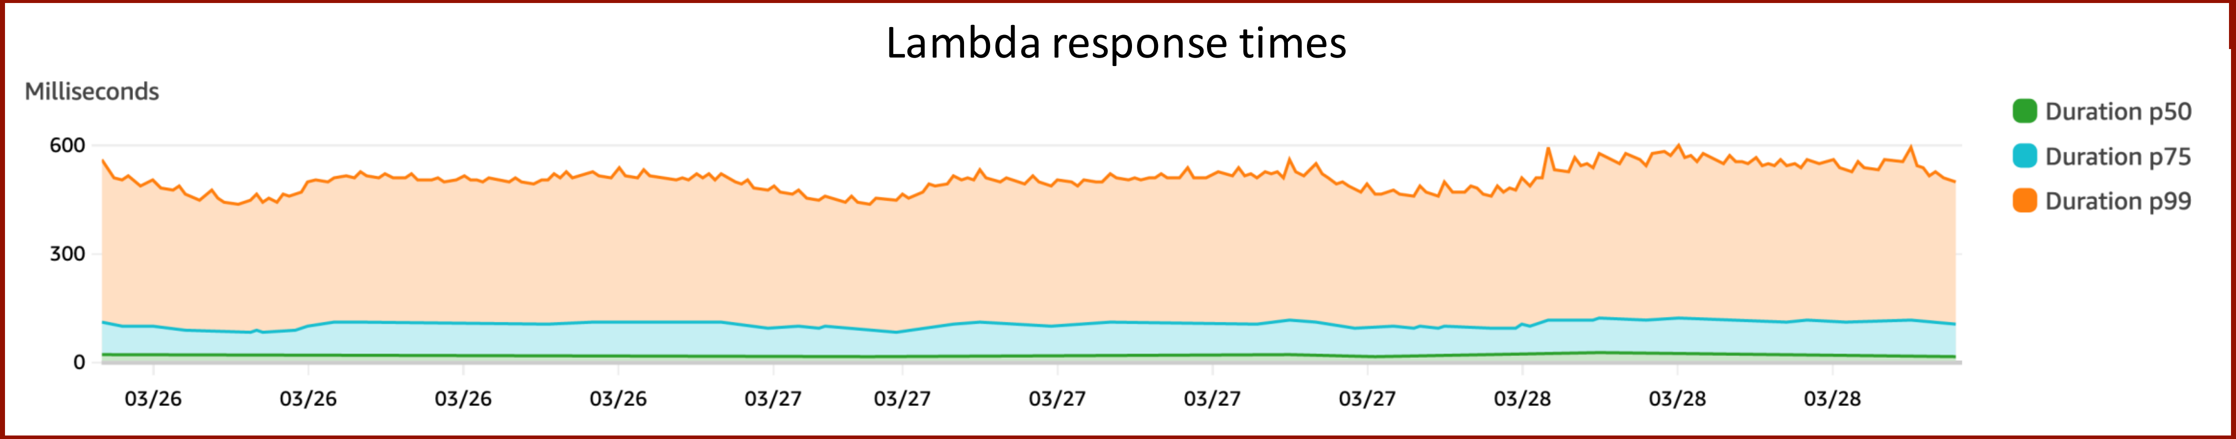

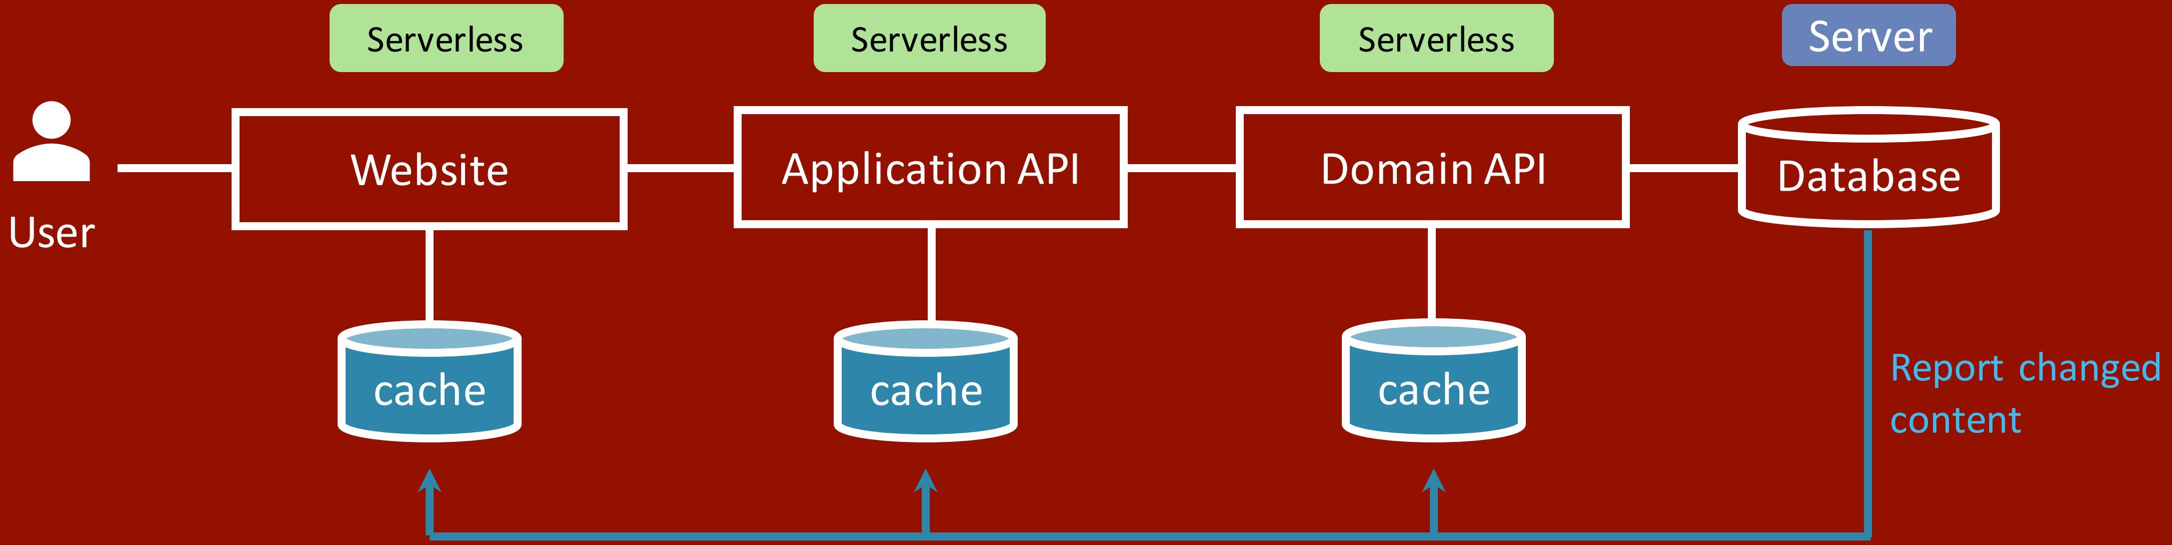

The page and the onward journeys are cached separately

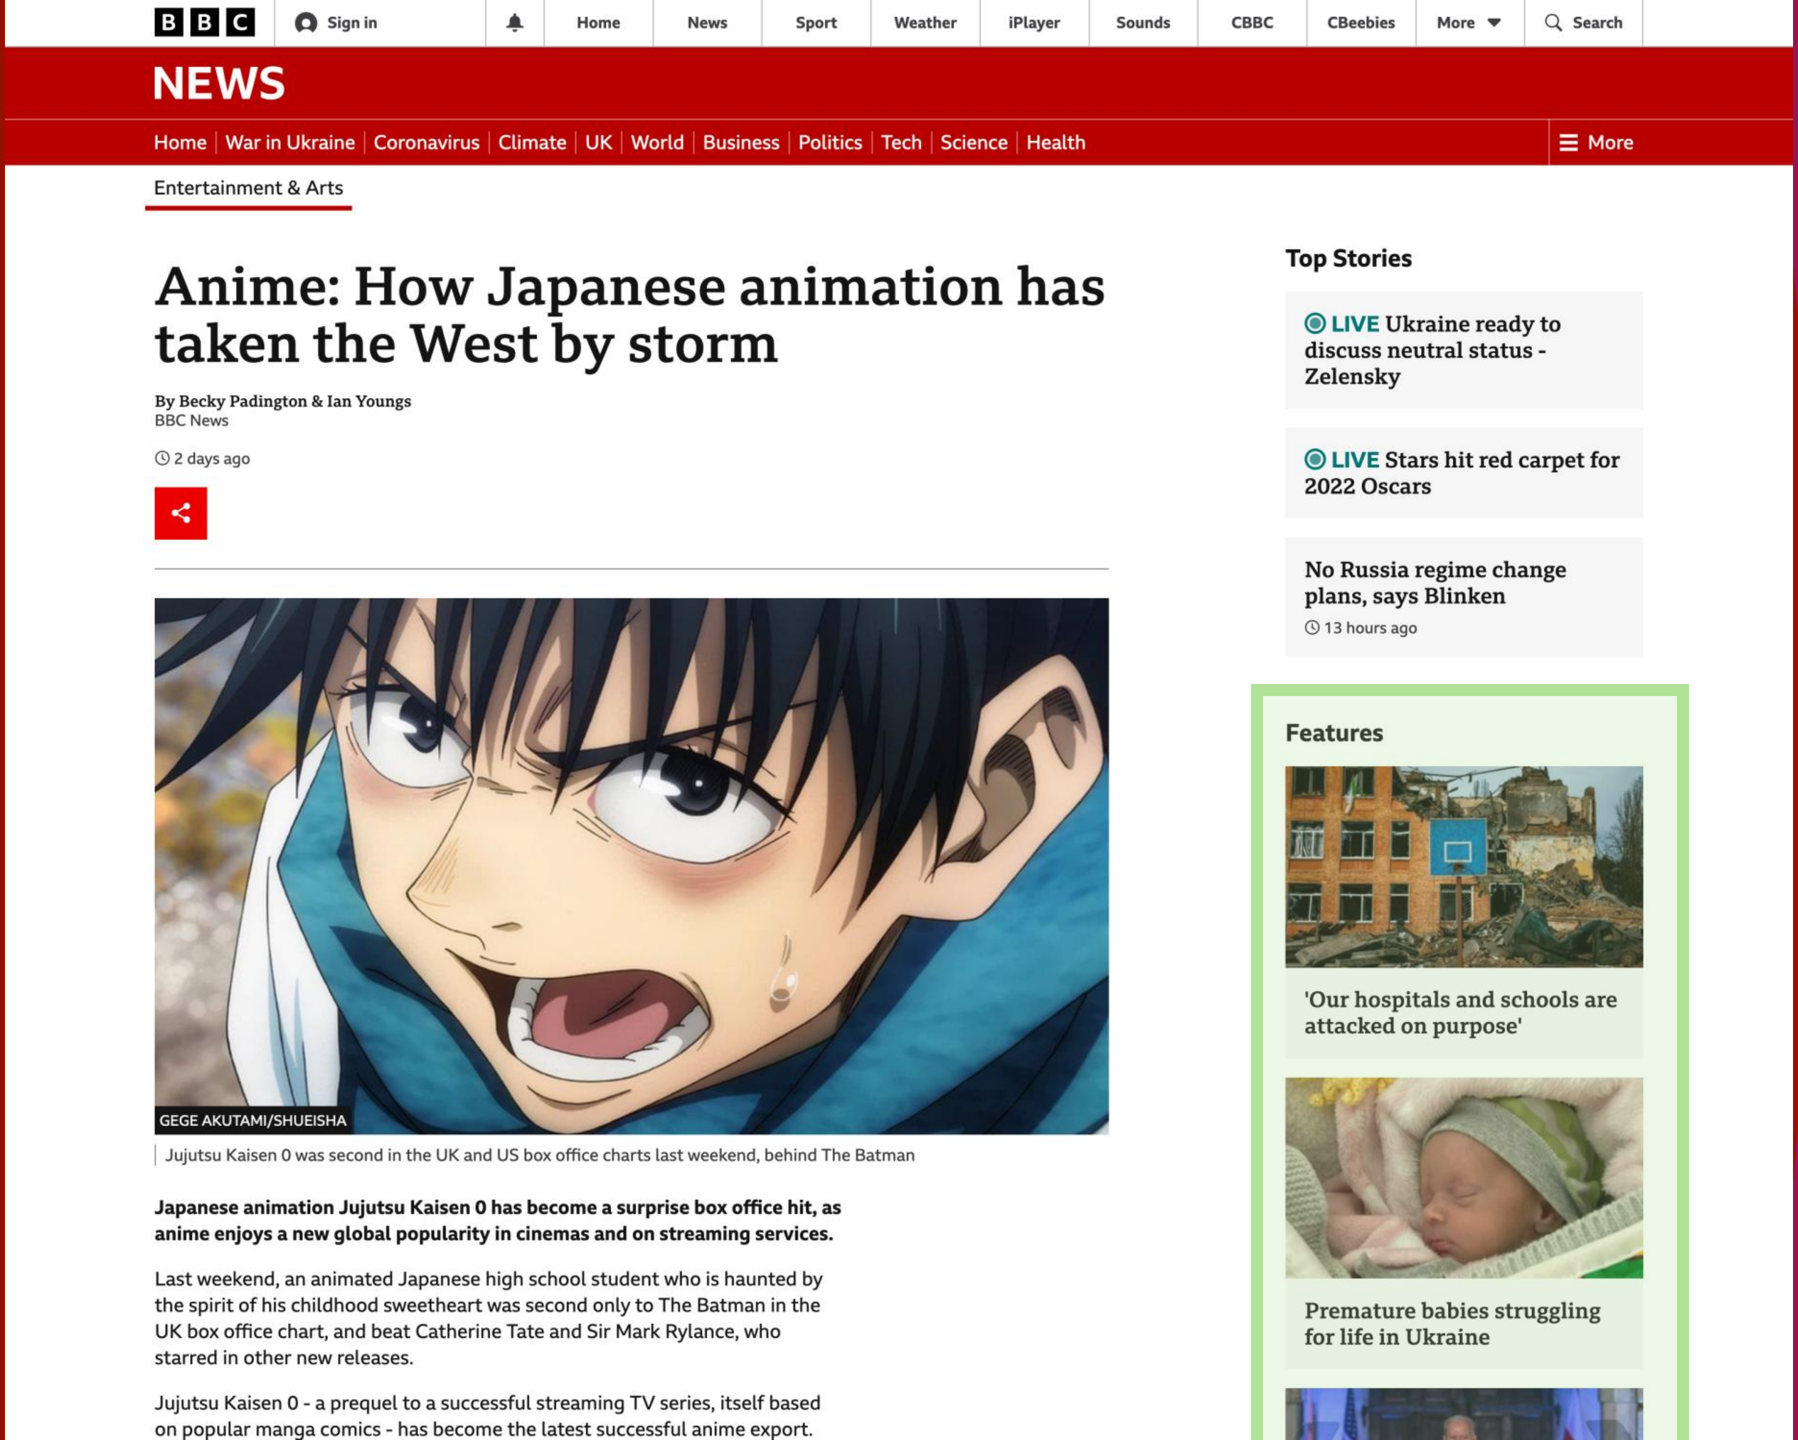

Anime: How Japanese animation has taken the West by storm

By Becky Padington & Ian Youngs
BBC News

2 days ago

Jujutsu Kaisen 0 was second in the UK and US box office charts last weekend, behind The Batman

Japanese animation Jujutsu Kaisen 0 has become a surprise box office hit, as anime enjoys a new global popularity in cinemas and on streaming services.

Last weekend, an animated Japanese high school student who is haunted by the spirit of his childhood sweetheart was second only to The Batman in the UK box office chart, and beat Catherine Tate and Sir Mark Rylance, who starred in other new releases.

Jujutsu Kaisen 0 - a prequel to a successful streaming TV series, itself based on popular manga comics - has become the latest successful anime export.

Top Stories

LIVE Ukraine ready to discuss neutral status - Zelensky

LIVE Stars hit red carpet for 2022 Oscars

No Russia regime change plans, says Blinken
13 hours ago

Features

'Our hospitals and schools are attacked on purpose'

Premature babies struggling for life in Ukraine

Manifest API

```
{
    "id" : "/news/123456",
    "lastUpdated" : 1648498767000,
    "title" : "Parrot learns seven times table"
},
{
    "id" : "/news/234567",
    "lastUpdated" : 1648496828000,
    "title" : ”Studies show fish happier on Fridays"
},
{
    "id" : "/news/345678",
    "lastUpdated" : 1648497939000,
    "title" : "Herd of elephants found in Norway"
}
```

```
{
    "id" : "/news/123456",
    "lastUpdated" : 1648498767000,
    "body" : "A parot in Italy has
learned to count in sevens. Billy, a
female macaw, regularly recites the
seven times table up to 105. Lorem
ipsum dolor sit amet, consectetur
adipiscing elit, sed do eiusmod
tempor incididunt ut labore et
dolore magna aliqua. Ut enim ad
minim veniam, quis nostrud
exercitation ullamco laboris nisi ut
aliquip ex ea commodo consequat.
Duis aute irure dolor in
reprehenderit in voluptate velit
esse cillum dolore eu fugiat nulla
pariatur. Excepteur sint occaecat
cupidatat non proident, sunt in
culpa qui officia deserunt mollit
anim id est laborum. Lorem ipsum
dolor sit amet, consectetur
adipiscing elit, sed do eiusmod
tempor incididunt ut labore et
dolore magna aliqua. Ut enim ad
```

One fast-updating, minimally cached API can provide references to another API with more immutable data

Designing APIs for **validation**

Schemas for the win!

- Provide a contract
- Enable stubbing by client teams
- Enable value adds to be built…

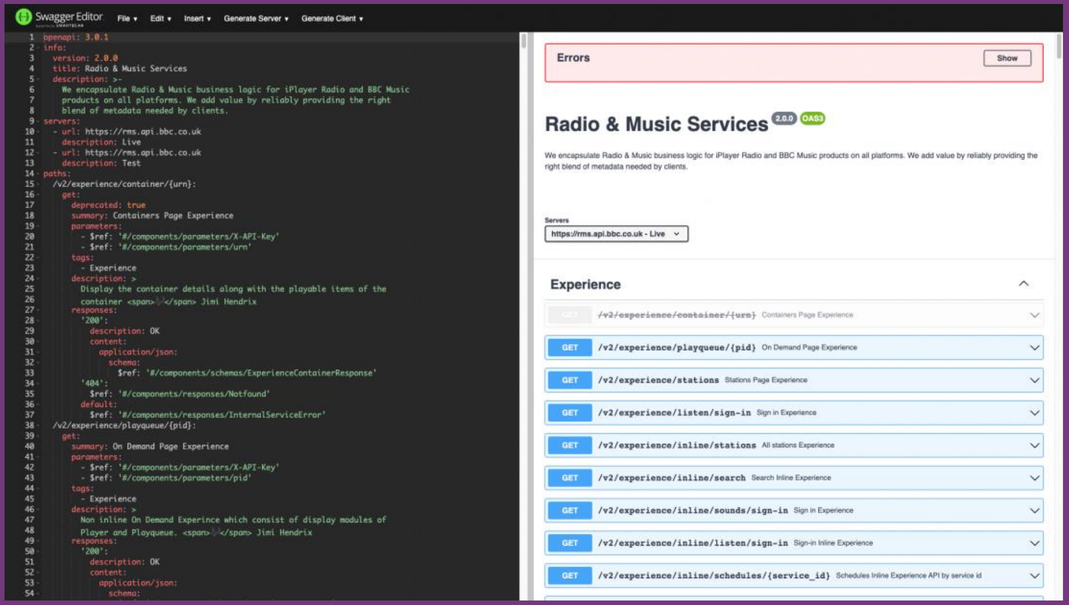

```yaml
openapi: 3.0.1
info:
  version: 2.0.0
  title: Radio & Music Services
  description: >-
    We encapsulate Radio & Music business logic for iPlayer Radio and BBC Music
    products on all platforms. We add value by reliably providing the right
    blend of metadata needed by clients.
servers:
  - url: https://rms.api.bbc.co.uk
    description: Live
  - url: https://rms.api.bbc.co.uk
    description: Test
paths:
  /v2/experience/container/{urn}:
    get:
      deprecated: true
      summary: Containers Page Experience
      parameters:
        - $ref: '#/components/parameters/X-API-Key'
        - $ref: '#/components/parameters/urn'
      tags:
        - Experience
      description: >
        Display the container details along with the playable items of the
        container <span></span> Jimi Hendrix
      responses:
        '200':
          description: OK
          content:
            application/json:
              schema:
                $ref: '#/components/schemas/ExperienceContainerResponse'
        '404':
          $ref: '#/components/responses/Notfound'
        default:
          $ref: '#/components/responses/InternalServiceError'
  /v2/experience/playqueue/{pid}:
    get:
      summary: On Demand Page Experience
      parameters:
        - $ref: '#/components/parameters/X-API-Key'
        - $ref: '#/components/parameters/pid'
      tags:
        - Experience
      description: >
        Non inline On Demand Experince which consist of display modules of
        Player and Playqueue. <span></span> Jimi Hendrix
      responses:
        '200':
          description: OK
          content:
            application/json:
              schema:
```

Errors

Radio & Music Services 2.0.0 OAS3

We encapsulate Radio & Music business logic for iPlayer Radio and BBC Music products on all platforms. We add value by reliably providing the right blend of metadata needed by clients.

Servers: https://rms.api.bbc.co.uk - Live

Experience

- GET ~~/v2/experience/container/{urn}~~ Containers Page Experience
- GET /v2/experience/playqueue/{pid} On Demand Page Experience
- GET /v2/experience/stations Stations Page Experience
- GET /v2/experience/listen/sign-in Sign in Experience
- GET /v2/experience/inline/stations All stations Experience
- GET /v2/experience/inline/search Search Inline Experience
- GET /v2/experience/inline/sounds/sign-in Sign in Experience
- GET /v2/experience/inline/listen/sign-in Sign-in Inline Experience
- GET /v2/experience/inline/schedules/{service_id} Schedules Inline Experience API by service id

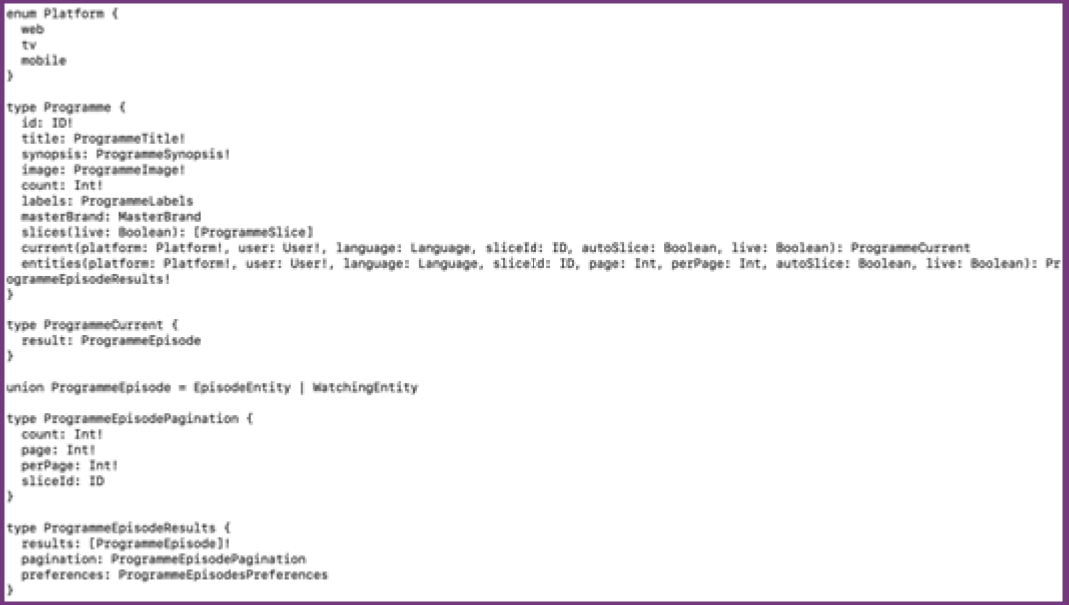

```graphql
enum Platform {
  web
  tv
  mobile
}

type Programme {
  id: ID!
  title: ProgrammeTitle!
  synopsis: ProgrammeSynopsis!
  image: ProgrammeImage!
  count: Int!
  labels: ProgrammeLabels
  masterBrand: MasterBrand
  slices(live: Boolean): [ProgrammeSlice]
  current(platform: Platform!, user: User!, language: Language, sliceId: ID, autoSlice: Boolean, live: Boolean): ProgrammeCurrent
  entities(platform: Platform!, user: User!, language: Language, sliceId: ID, page: Int, perPage: Int, autoSlice: Boolean, live: Boolean): Pr
ogrammeEpisodeResults!
}

type ProgrammeCurrent {
  result: ProgrammeEpisode
}

union ProgrammeEpisode = EpisodeEntity | WatchingEntity

type ProgrammeEpisodePagination {
  count: Int!
  page: Int!
  perPage: Int!
  sliceId: ID
}

type ProgrammeEpisodeResults {
  results: [ProgrammeEpisode]!
  pagination: ProgrammeEpisodePagination
  preferences: ProgrammeEpisodesPreferences
}
```

Unreliable Upstream

Like your upstream services but worse…

- Data that meets their API
- Random response times and occasional errors
- Response profiles with configured delays and error rates.

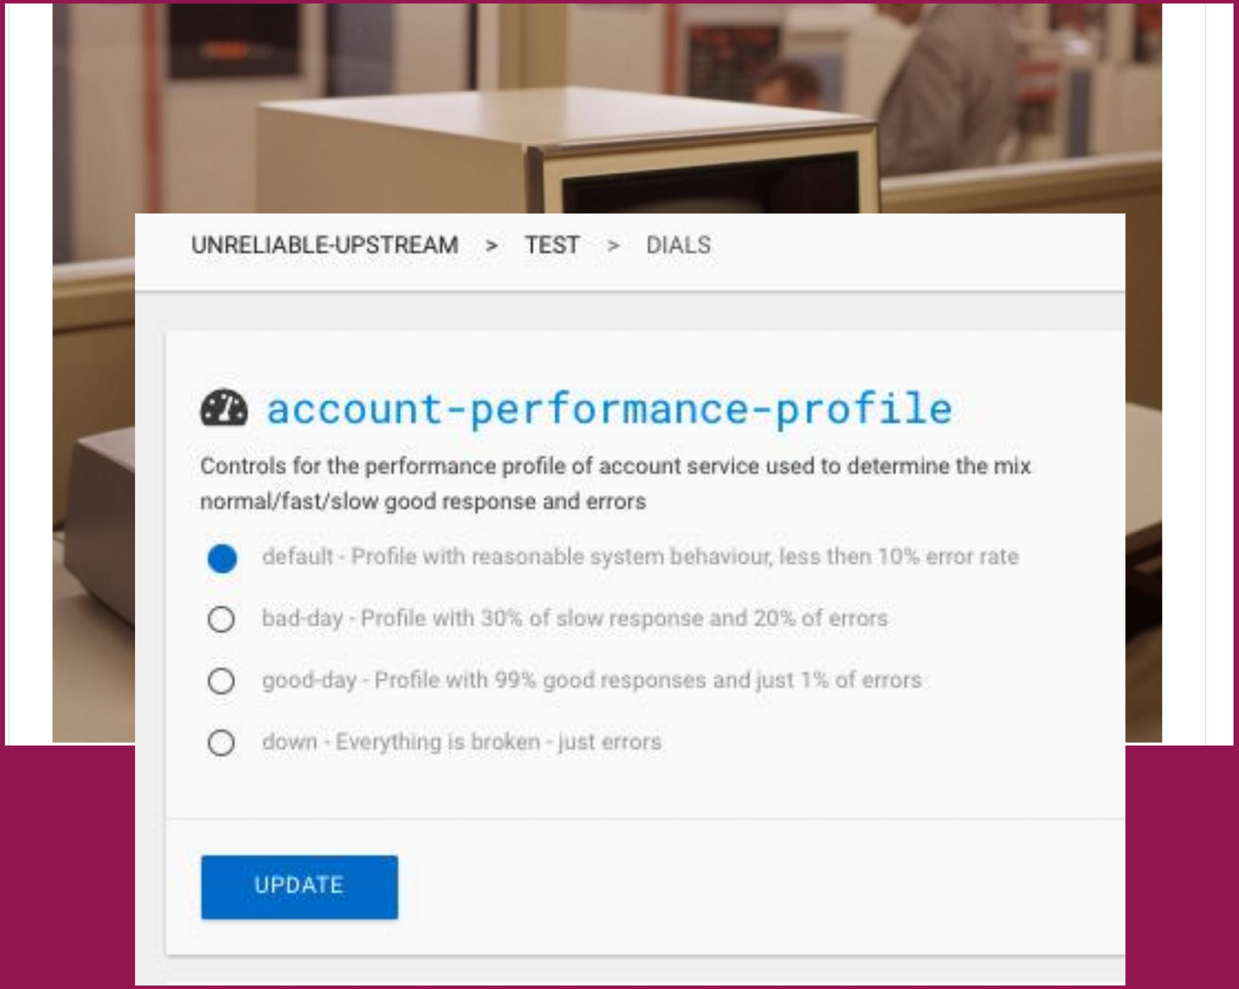

UNRELIABLE-UPSTREAM > TEST > DIALS

account-performance-profile

Controls for the performance profile of account service used to determine the mix normal/fast/slow good response and errors

● default - Profile with reasonable system behaviour, less then 10% error rate

○ bad-day - Profile with 30% of slow response and 20% of errors

○ good-day - Profile with 99% good responses and just 1% of errors

○ down - Everything is broken - just errors

Designing for **frontends**

Designing for **change**

Designing for **scale**

Designing for **validation**

Summary

Focus on doing one thing well

Know your clients

Serverless helps APIs with speed and scale

Treat failure as a first-class citizen

Matthew Clark

@matthew1000

Paul Caporn

@drpacman72

BBC