

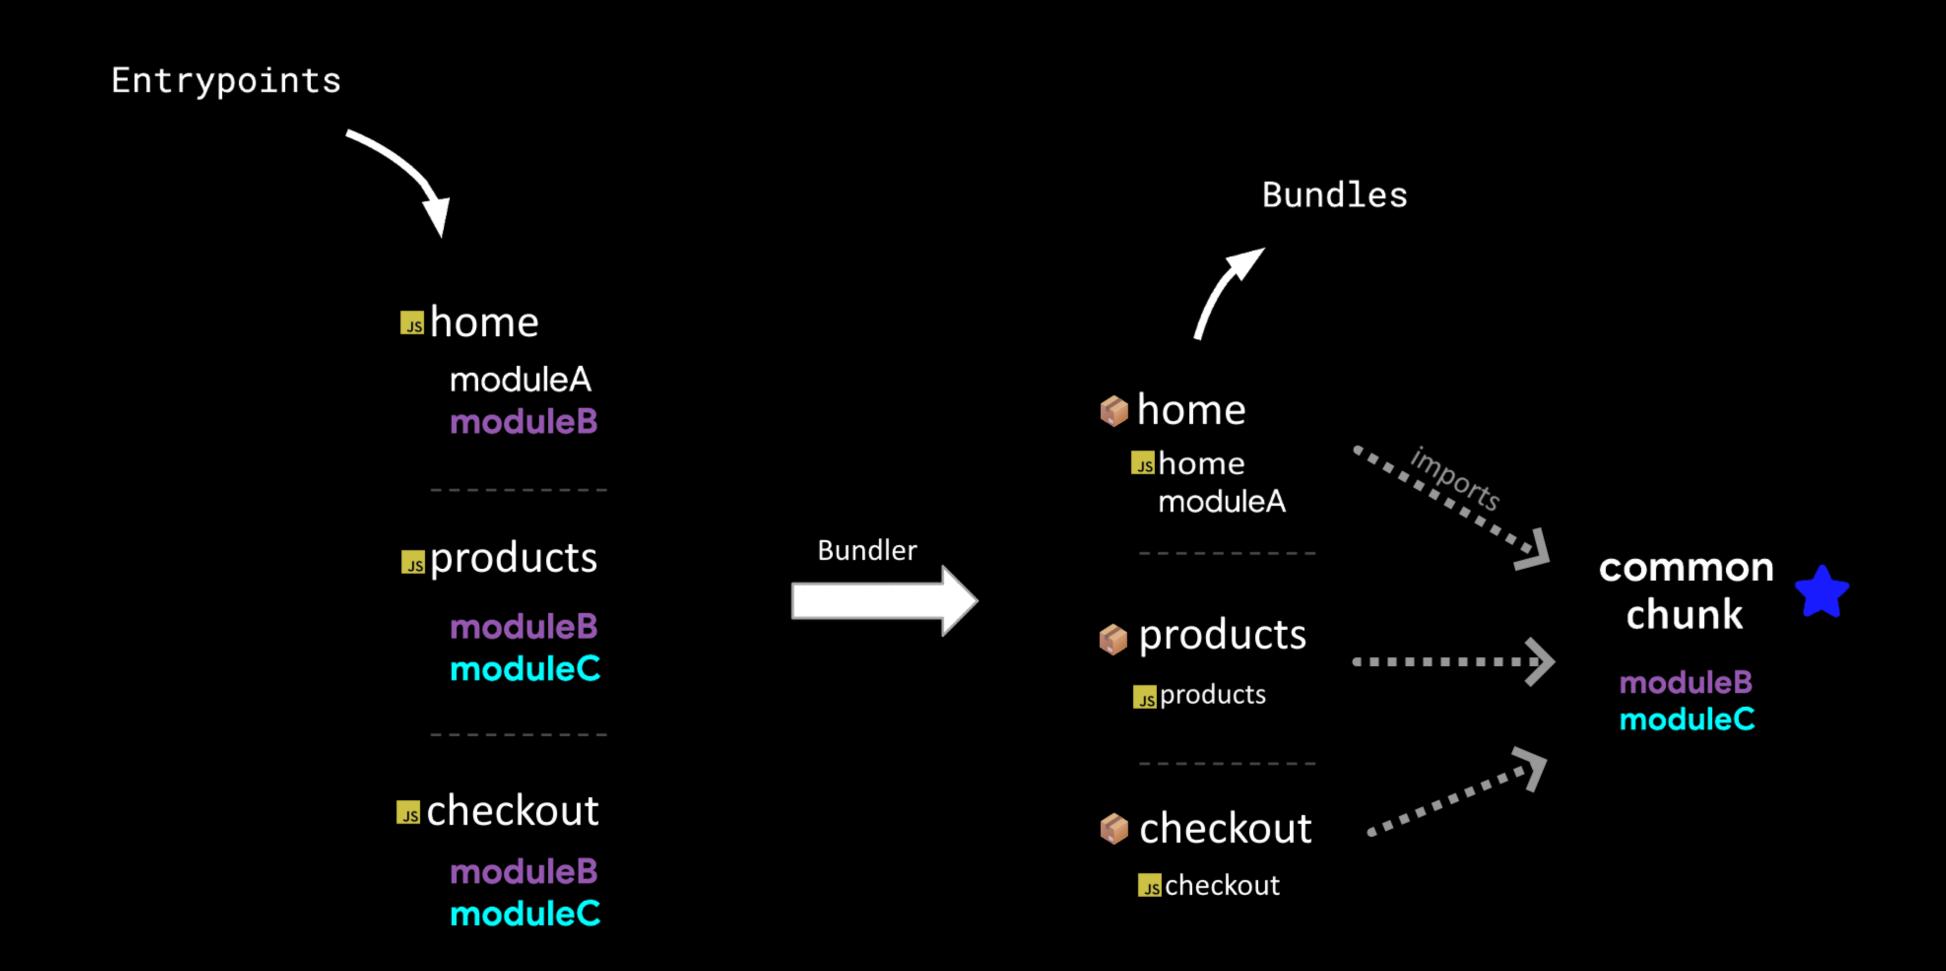
## Rolldown: Chunking in the wild

#### Who am !?

- Senior Performance Engineer @ Framer
- I work on Core Web Vitals & performance related topics
- Participant in the W3C WebPerf Working Group come chat w/ me about perf topics :)

Framer is a website builder and design tool, where you can design any website for any scale, setup complex animations, user interactions, ... and have it published in seconds.

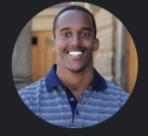
### What are "chunks"?



## How many chunks are optimal?

# Improved Next.js and Gatsby page load performance with granular chunking

A newer webpack chunking strategy in Next.js and Gatsby minimizes duplicate code to improve page load performance.



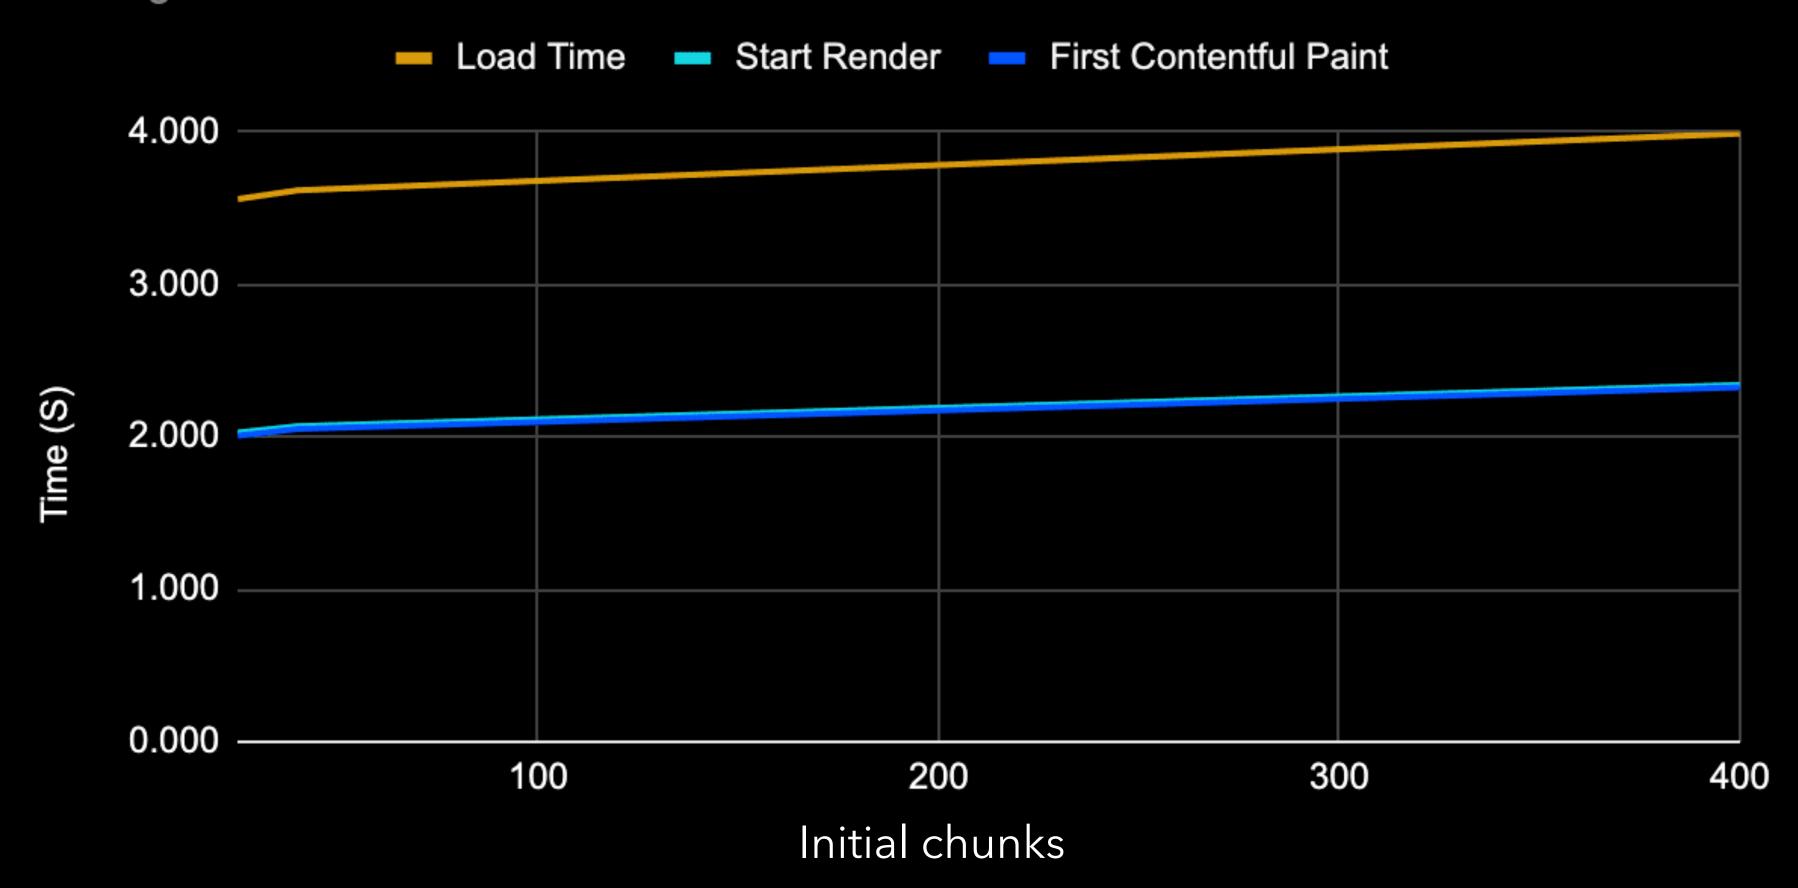
Houssein Djirdeh



Chrome is collaborating with tooling and frameworks in the JavaScript open-source ecosystem. A number of newer optimizations were recently added to improve the loading performance of Next.js and Gatsby. This article covers an improved granular chunking strategy that is now shipped by default in both frameworks.

## How many chunks are optimal?

Average of three runs on Zeit.co



# 20-25

"Initial chunks struck the right balance between loading performance and caching efficiency"

This talk isn't meant to bash on esbuild.

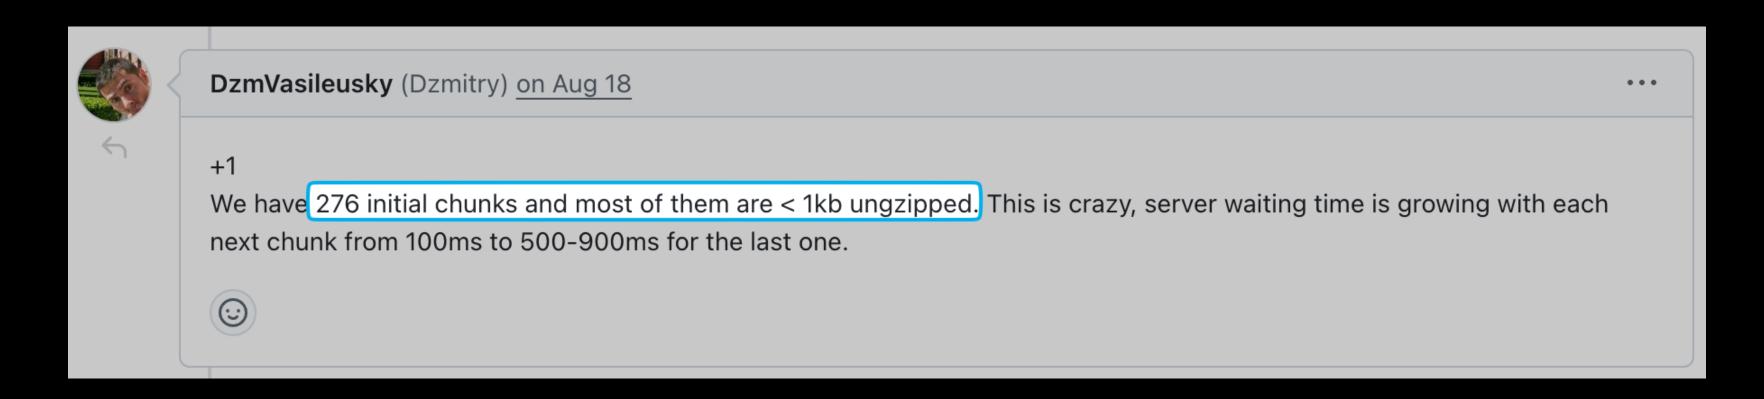
esbuild served us great. We were (one of?) the first users of esbuild in production.

It's stable & incredibly fast — Evan Wallace started the race of fast bundlers.

But: It lacks some options – a void for other bundlers to fill.

## Too many small chunks





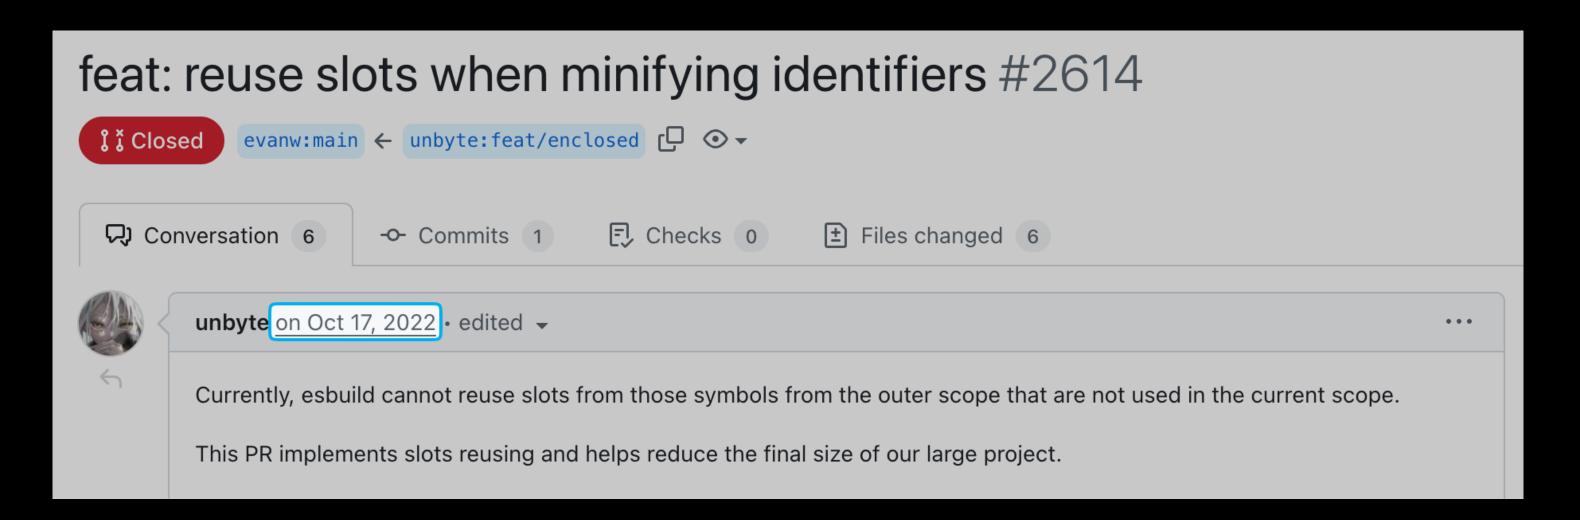
## Not so strict import order

Incorrect import order with code splitting and multiple entry points Open iamakulov (Ivan Akulov) opened n Sep 20, 2020 · dited by iamakulov For the workaround, see #399 (comment). Consider the following code: // entry1.js import "./init-dep-1.js"; import "./run-dep.js"; // entry2.js import "./init-dep-2.js"; import "./run-dep.js"; // init-dep-1.js global.foo = { --log: () => console.log("foo.log() (from entry 1) called"), // init-dep-2.js global.foo = { ...log: () => console.log("foo.log() (from entry 2) called"), // run-dep.js global.foo.log(); When you bundle this code with esbuild ——bundle ——outdir=build ——format=esm ——splitting ——platform=node ——out extension:.js=.mjs ./entry1.js ./entry2.js , ESBuild discovers that run-dep.js is a module shared between both entry points. Because code splitting is enabled, ESBuild moves it into a separate chunk: O // build/entry1.js import "./chunk.P2RPFHF3.js"; global.foo = {

However, by doing so, ESBuild puts run-dep.js above the init-dep-1.js or init-dep-2.js code. This changes the import

..log: () => console.log("foo.log() (from entry 1) called")

### Stale PRs



Artifact	Original size	Gzip size	
react v17.0.2 (Source)	72.13 · kB	19.40 · kB	
Minifier	Minified size	Minzipped size	Time
SWC	<b>№</b> -68% 22.89 kB	<b>№</b> -58% <b>8.18</b> kB	<b>22</b> ⋅ ms
<u>terser</u>	-68% 23.14 kB	-57% 8.32 kB	<sup>7x</sup> 153 · ms
esbuild-optimized	-67% 23.63 kB	-56% 8.51 kB	<sup>1</sup> x 22 · ms
<u>esbuild</u>	-67% 23.70 kB	-56% 8.53 kB	<sup>1</sup> x 22 ms

Alipay internally forked esbuild

### Fork it?

## Unbundle everything?





# Why bundle in 2025

## Why bundle?

#### HTTP/2+ does not solve the need for bundling.

Effects of splitting one file into multiple ones:

Network overhead: 1 roundtrips, 1 compression, HTTP header overhead



As a rule, RTT ( $\alpha$ ) stays constant while download time (x) is proportional to filesize. Therefore, splitting one large bundle into 16 smaller ones goes from 1 $\alpha$  + x to 16 $\alpha$  + 16(0.0625x). Expect things to probably get a little slower.

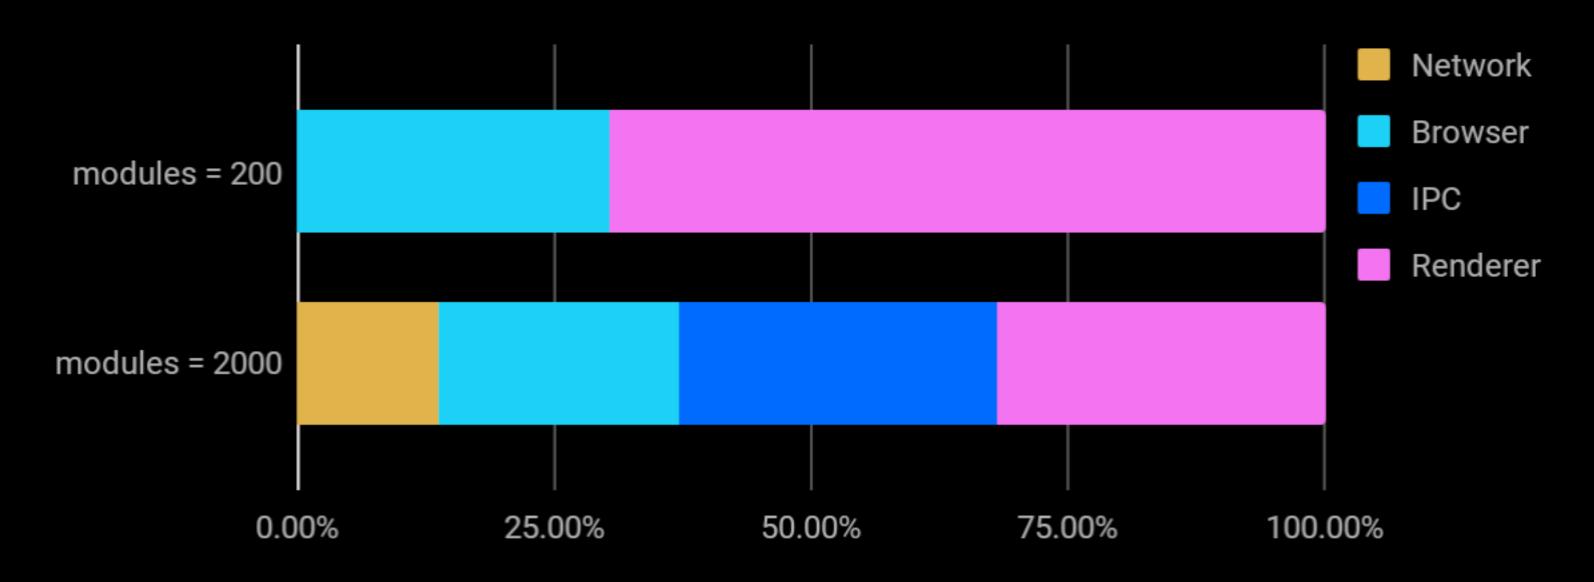


HTTP/2+ doesn't enable us to extend a compression window beyond a single resource. What that means in practice is that the compression ratio of a lot of small files is significantly worse than the compression ratio of fewer, larger files.

Source: <u>Harry Roberts</u>

### Renderer & IPC cost

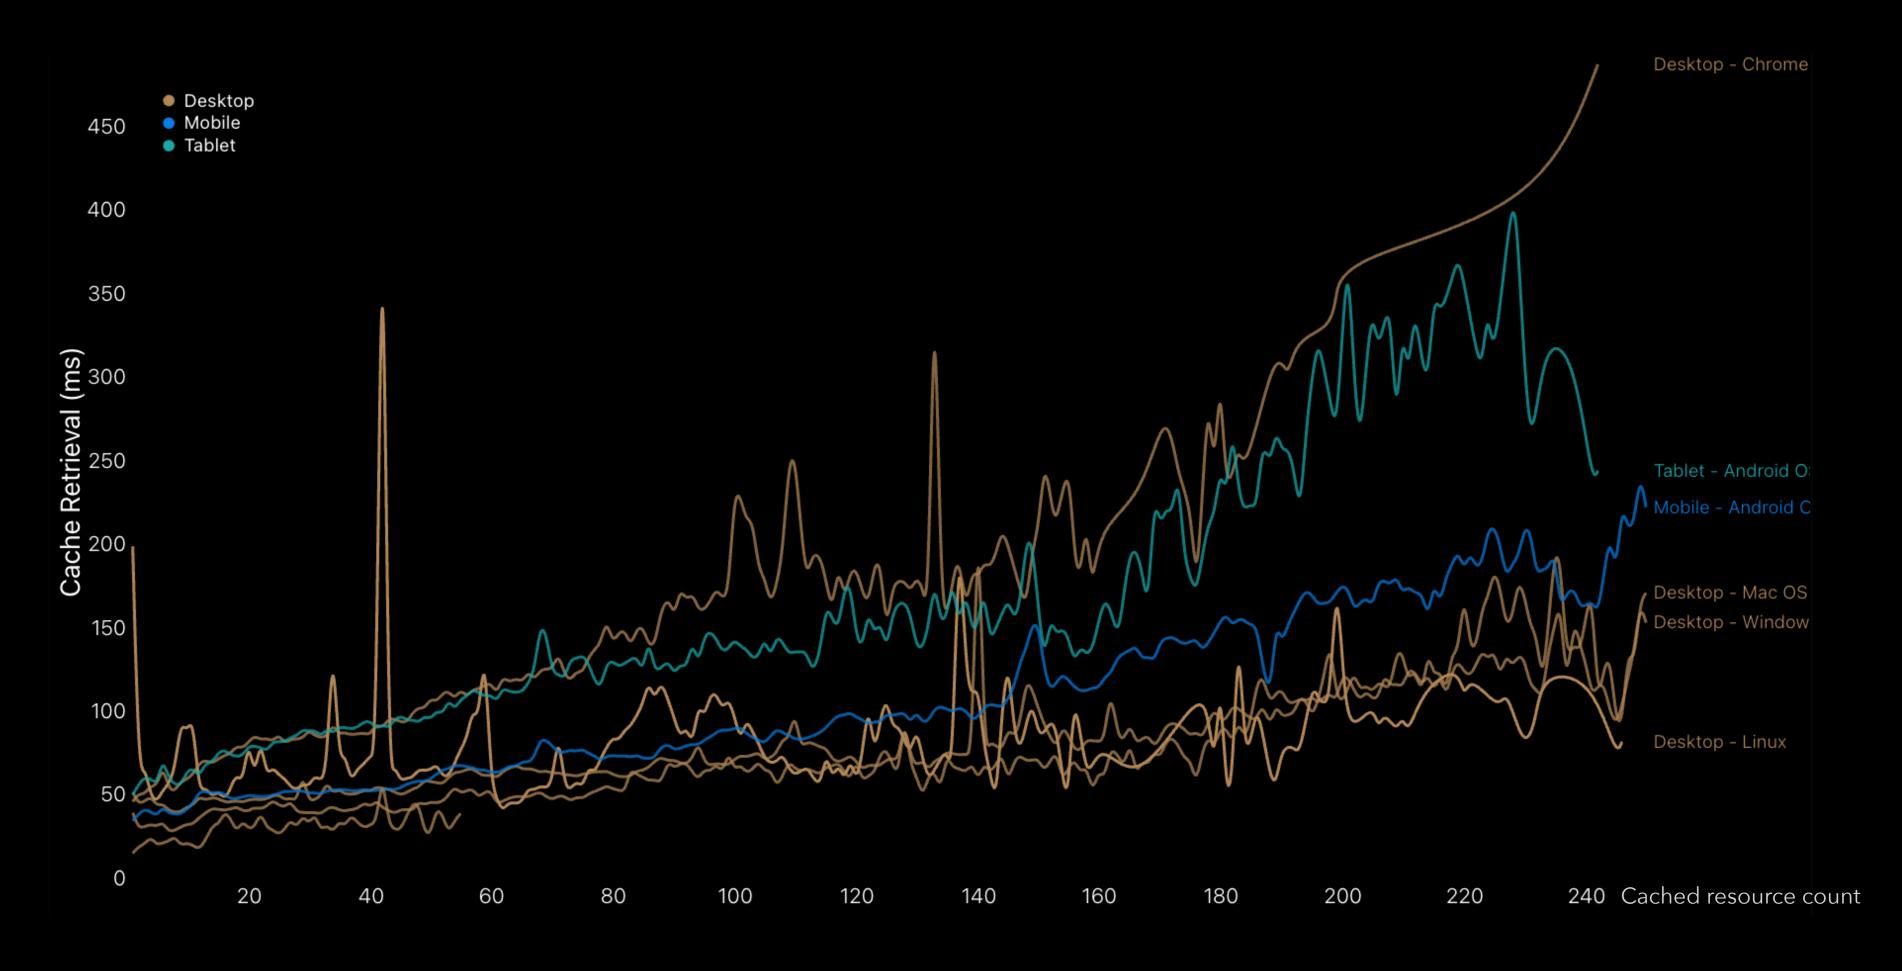
Modules have a cost in both Blink, V8 and the communication between.



"Loading time increases proportionally to the # of requests, not by the total size of responses, regardless of whether they come from cache or not. For example:

Loading 10\_000 1kB files is 10x slower than a single 10MB load."

### Renderer & IPC cost



Chrome OS average cache retrieval doubles from  $\pm 50$ ms with 5 cached resources up to  $\pm 100$ ms with 25 resources.

#### Geekbench 5 Multi-Core Scores iPhone 15 Pro ■ iOS Multi-Core ■ Android Multi-Core ■ Budget Android Multi-Core ■ Low-End Multi-Core iPhone 149Pro Galaxy S23 Ultra 5000 iPhone 1 iPhone 12 Galaxy S21 Ultr@alaxy S22 Ultra iPhone 1 Galaxy S20 Ultra iPhone XS Galaxy S10+ 2500 Galaxy S9+ iPhone 8 Galaxy A24 (4G) Galaxy S8+ iPhone 7 Galaxy S7 Galaxy A50 Galaxy A51 Galaxy A51 Galaxy A23 (40) Moto G7 iPhone 6s Honor 7X Moto E7 Moto E30 Moto E Galaxy S6 Gallaxy S5 Galaxy S4 Moto G4+ Moto G4 Moto E6 Nokia 2.1 Nokia 2 Moto E Sony Xperia M 2013 2014 2015 2016 2017 2018 2019 2020 2021 2022 2023

### Think about all users

## Concurrency limits

HTTP/2+ servers only accept **N concurrent streams**. After, they kinda behave like HTTP/1.

Table 6.1 Concurrent stream limits on popular HTTP/2 server-side implementations

Software	Туре	Default concurrent streams	
Apache HTTPD (v2.4.35)	Web server	100 <sup>a</sup>	
nginx (v1.14.0)	Web server	128 <sup>b</sup>	
H20 (2.3.0)	Web server	100°	
IIS (v10)	Web server	100	
Jetty (9.4.12)	Web and Java servlet container	128 <sup>d</sup>	
Apache Tomcat (9.0)	Web and Java servlet container	200 <sup>e</sup>	
Node (10.11.0)	JavaScript runtime environment	100 <sup>f</sup>	
Akamai	CDN	100	
Amazon CloudFront and S3	CDN	128	
Cloudflare	CDN	128	
MaxCDN	CDN	128	

Table 6.2 Concurrent stream limits on popular HTTP/2 web browsers

Software	Default concurrent streams	
Chrome (v69)	1000	
Firefox (v62)	Not set (uses HTTP/2 default of unlimited)	
Safari (v12)	1000	
Opera (v56)	1000	
Edge (v17)	1024	
Internet Explorer 11	1024	

#### You should bundle.

#### HTTP/2+ does not solve the need for bundling.

Splitting one file into multiple ones creates more overhead:

- 1. Network overhead: † roundtrips, ↓ compression, HTTP header overhead
- 2. Renderer & IPC cost: More isn't free.
- 3. Concurrency limits of HTTP/2+ servers & browsers

This leads to worse UX & Core Web Vitals – especially on slower devices & networks.



## Ambitious bundler goals

Our 3 must-have requirements:

- 1. Granular control over chunks to counter too-small-modules.
- 2. Strict module execution order to isolate user code in our case.
- 3. Equally or faster than esbuild regressing perf isn't an option.

## A bundler journey

There are other good bundlers out there:

- Parcel: Surprisingly fast and capable. Slower than esbuild.
- Rspack: Similar API to Webpack. In my tests in 2024, slower than Parcel.
- Farm-Fe: Unique chunking algo. Slower than Rspack.
- Turbopack: Tied to Next.js.





Back in May 2024, we knew Rolldown was coming, but it was early.

We've seen how fast Oxc is. We heard Rolldown is fast.

## Rolldown at 7 Framer

from research to product

After talking in August 2024, we figured a cooperation would shape Framer's future of bundling and I could help shape Rolldown's API.

## What happened in 1 year?

- Designing & implementing the chunking API
- Minification support
- Fix rare minification errors
- Handle legal comments during minification
- Fix crashes
- Hashing woes
- Top-level await support
- CJS in ESM
- Nasty chunking algorithm things (Yunfei can tell you a story)
- Performance & memory improvements (bundle 300 MB JS?)
- Test Rolldown for Framer sites with 1 route and with 10\_000s of routes

## Esbuild vs Rolldown

(from my POV)

- Similar API (with unique Rolldown additions)
- JS plugin support (we use it to create virtual modules)
- JS minification
- X No CSS minification yet!

## Minifier performance

Artifact	Original size	Gzip size	
typescript v4.9.5 (Source)	10.95 MB	1.88 MB	
Minifier	Minified size	Minzipped size	Time
1. oxc-minify	-70% 3.34 MB	<b>₹-55%</b> 855.08 KB	<sup>5x</sup> 607 ms
2. <u>@tdewolff/minify</u>	-69% 3.35 MB	<sup>-54%</sup> 875.82 KB	<sup>2x</sup> 255 ms
3. @swc/core	<b>₹-70%</b> 3.31 MB	<sup>-54%</sup> 859.05 KB	<sup>13x</sup> 1,487 ms
4. <u>esbuild</u>	<sup>-68%</sup> 3.49 MB	<sup>-51%</sup> 915.59 KB	<sup>4x</sup> 485 ms
5. <u>bun</u>	<sup>-68%</sup> 3.54 MB	<sup>-51%</sup> 923.27 KB	<sup>2x</sup> 293 ms
6. <u>uglify-js (no compress)</u>	-68% 3.54 MB	<sup>-53%</sup> 876.54 KB	<sup>34x</sup> 3,784 ms
7. terser (no compress)	-68% 3.53 MB	<sup>-53%</sup> 878.64 KB	<sup>43x</sup> 4,687 ms

#### Esbuild vs Rolldown

Memory usage when bundling 300 MB JS w/ 500 MB sourcemaps:

Rolldown: >10 GB

esbuild: 9.2 GB

Improvements are coming! (And I hope no one else has 300 MB of JS somewhere)

## Bundling at Framer

 Customers can write custom JS code, most commonly React components → our biggest sites have ±300 MB of JS (across all pages)

We use SSG for faster page loads & crawl-ability (no SSR → no LLM!)

- Rolldown bundles both client and server code
  - → 🔊 Deps needed for SSR shouldn't end up in the client bundle
  - → Strict execution order matters our runtime code shouldn't fail if there is an error in user code

## Bundling at Framer

- Every route has its own entry-point
- → Lazy loaded in the browser (& preloaded using heuristics)

- Every route might have its own dependencies
- → They shouldn't be loaded on every other route

Non trivial: How do you split the code optimally?

## Good rules for bundling

#### Google research to the rescue:

- A. Used by some chunks and below 20 kb
  - → "common" chunks
- B. Used on all pages
  - → "shared" chunks
- C. Big library > 160kb
  - → "library" chunks

## Start simple

```
groups:
    name: "react",
    // include react/react-dom/scheduler, but not the server chunk
    test: /[\\/]npm:(react|react-dom|scheduler)@(?![\d.]+\/server\.browser\.js)/,
     priority: 1000,
   name: "motion", test: /^(motion|framer-motion)/, },
                                                                 Update independently
   name: "framer", test: /^framer/, },
    name: "shared-lib",
                                                                 Used by all routes
    minShareCount: routesLength,
  routesLength >= 40 ? {
                                                                 More likely to hit
    name: "shared",
     minShareCount: Math.floor(routesLength / 2),
                                                                 20-25 chunks
  }: undefined,
].filter(group => group !== undefined),
```

## Advanced Options in Rolldown

Goal: Simple, yet powerful options similar to webpack.

- A. minSize → don't create the group unless > N bytes
- B. maxSize → split the group when > N bytes
- C. minShareCount
  - create only if N entrypoints reference it
- D. min/maxModuleSize
  - → min or max size of modules that are bundled into the group

#### How Framer loads JS

Browser preload scanner starts downloading the chunks before executing "main"

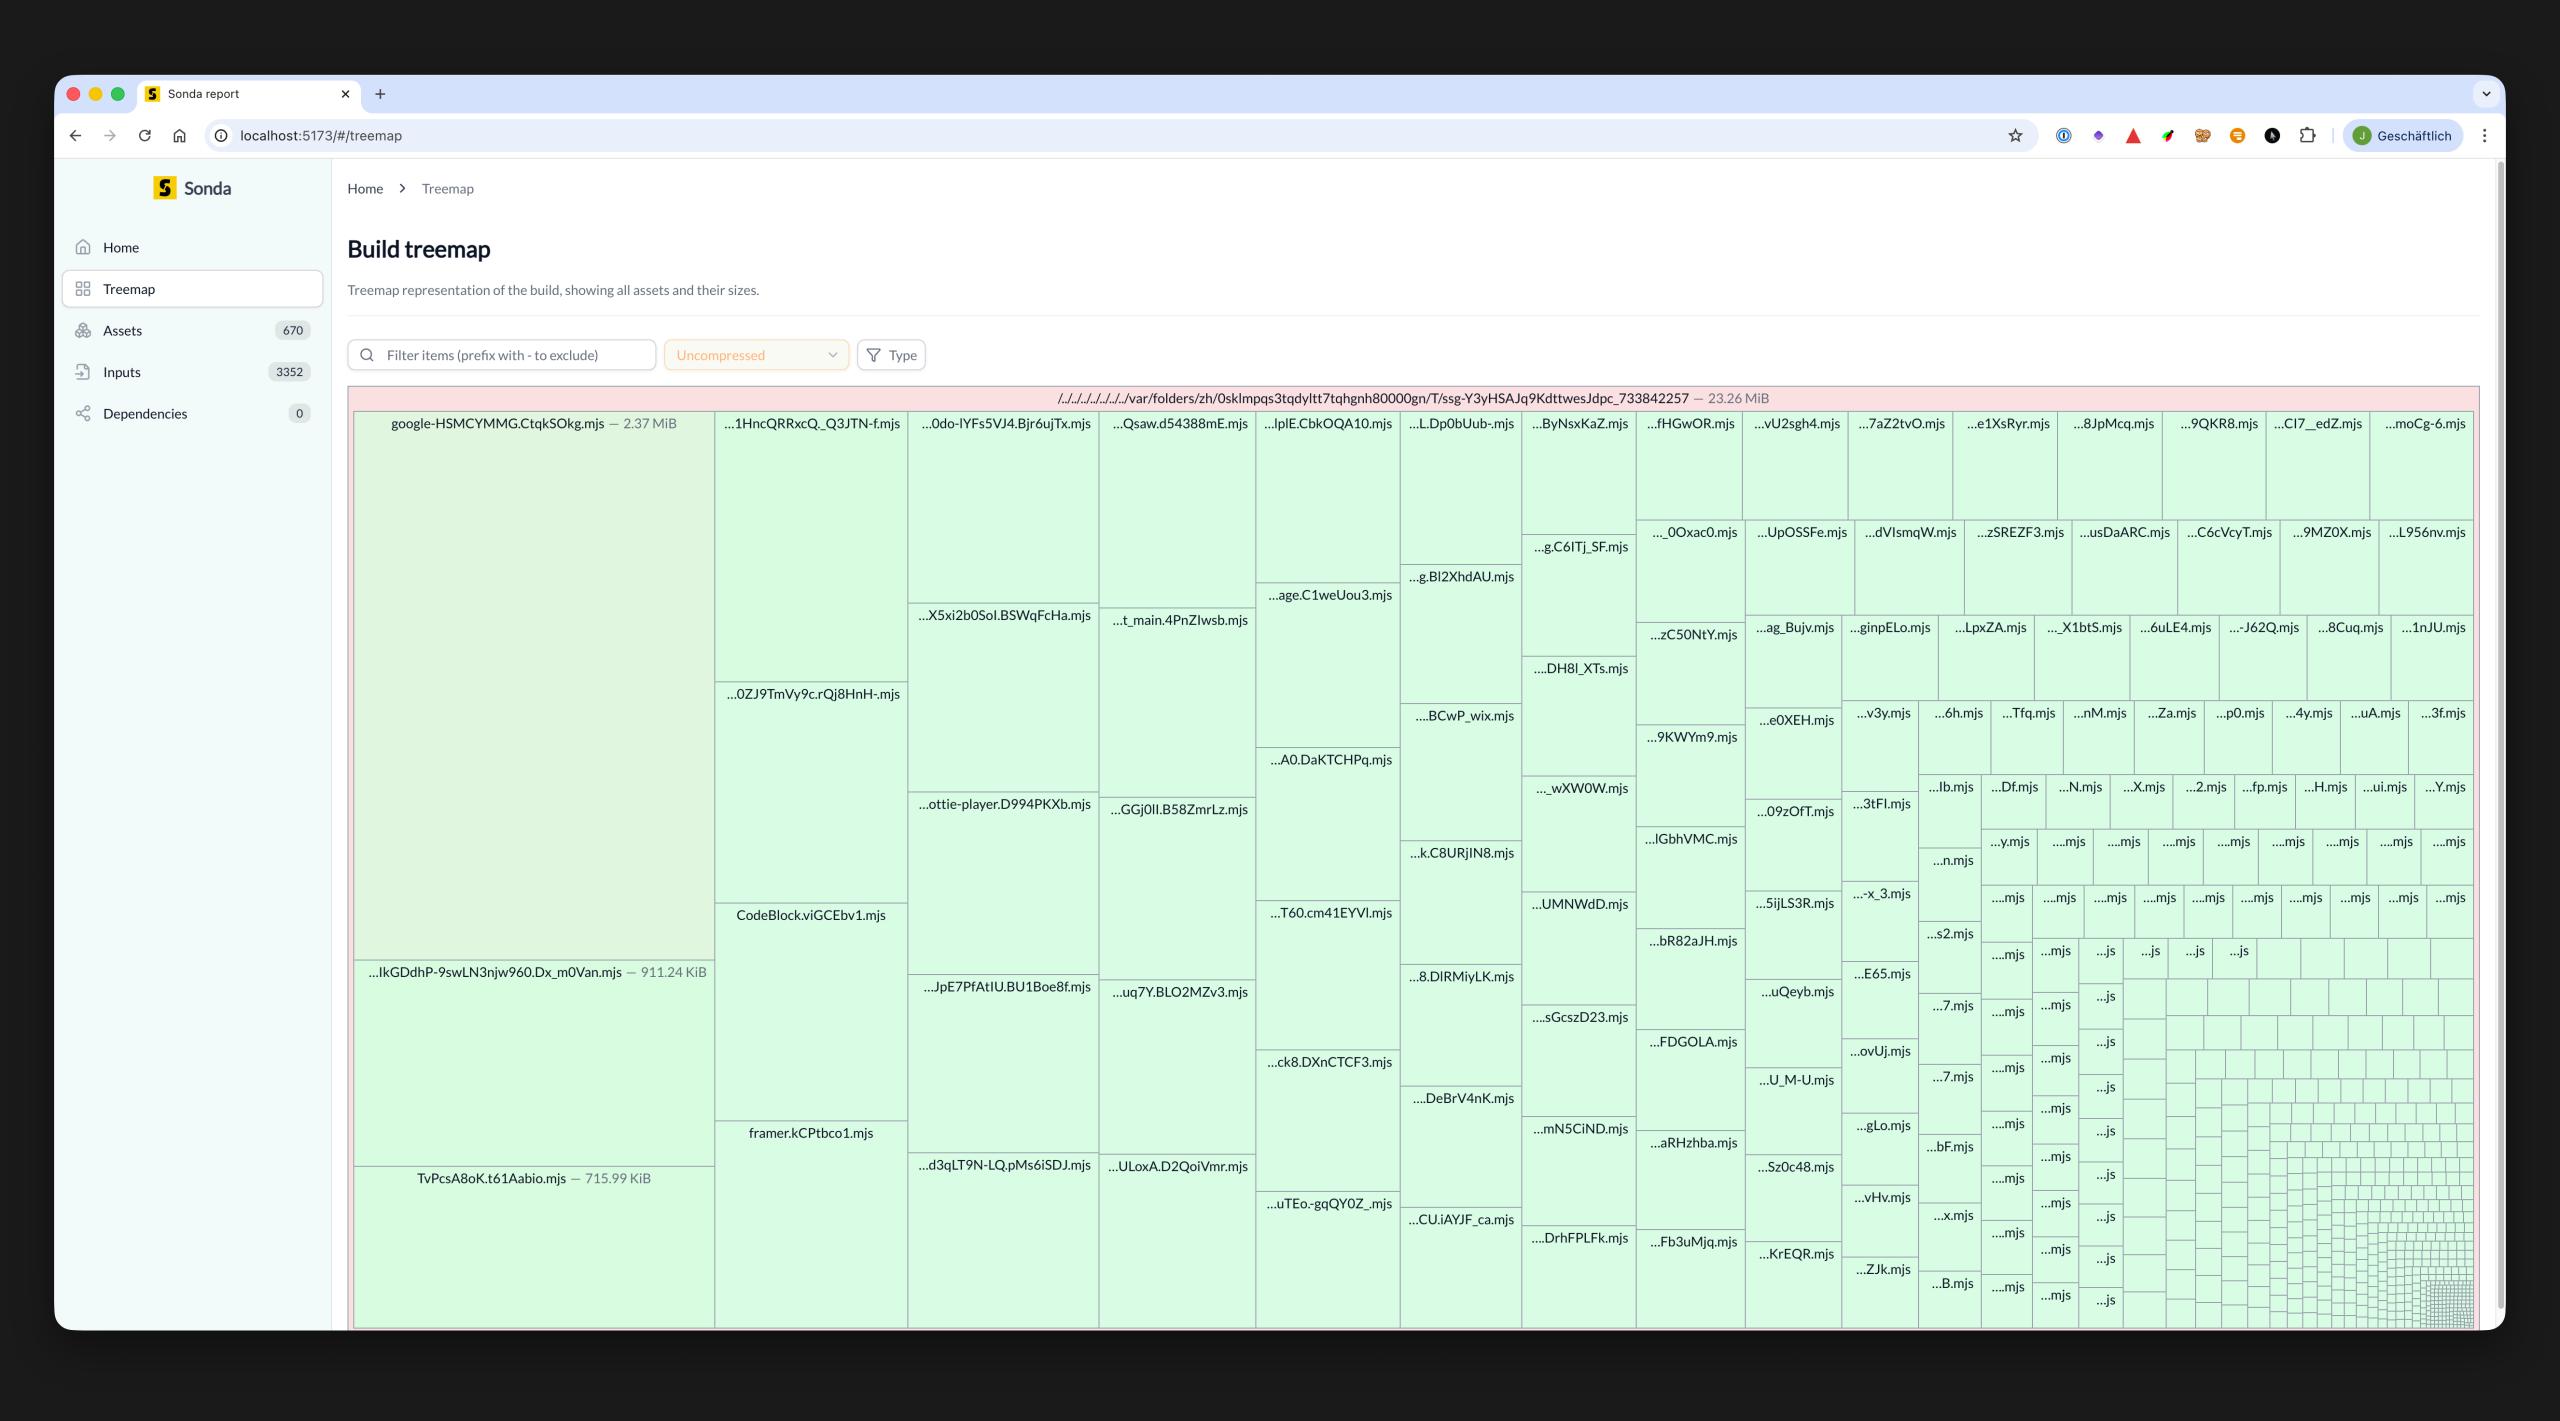


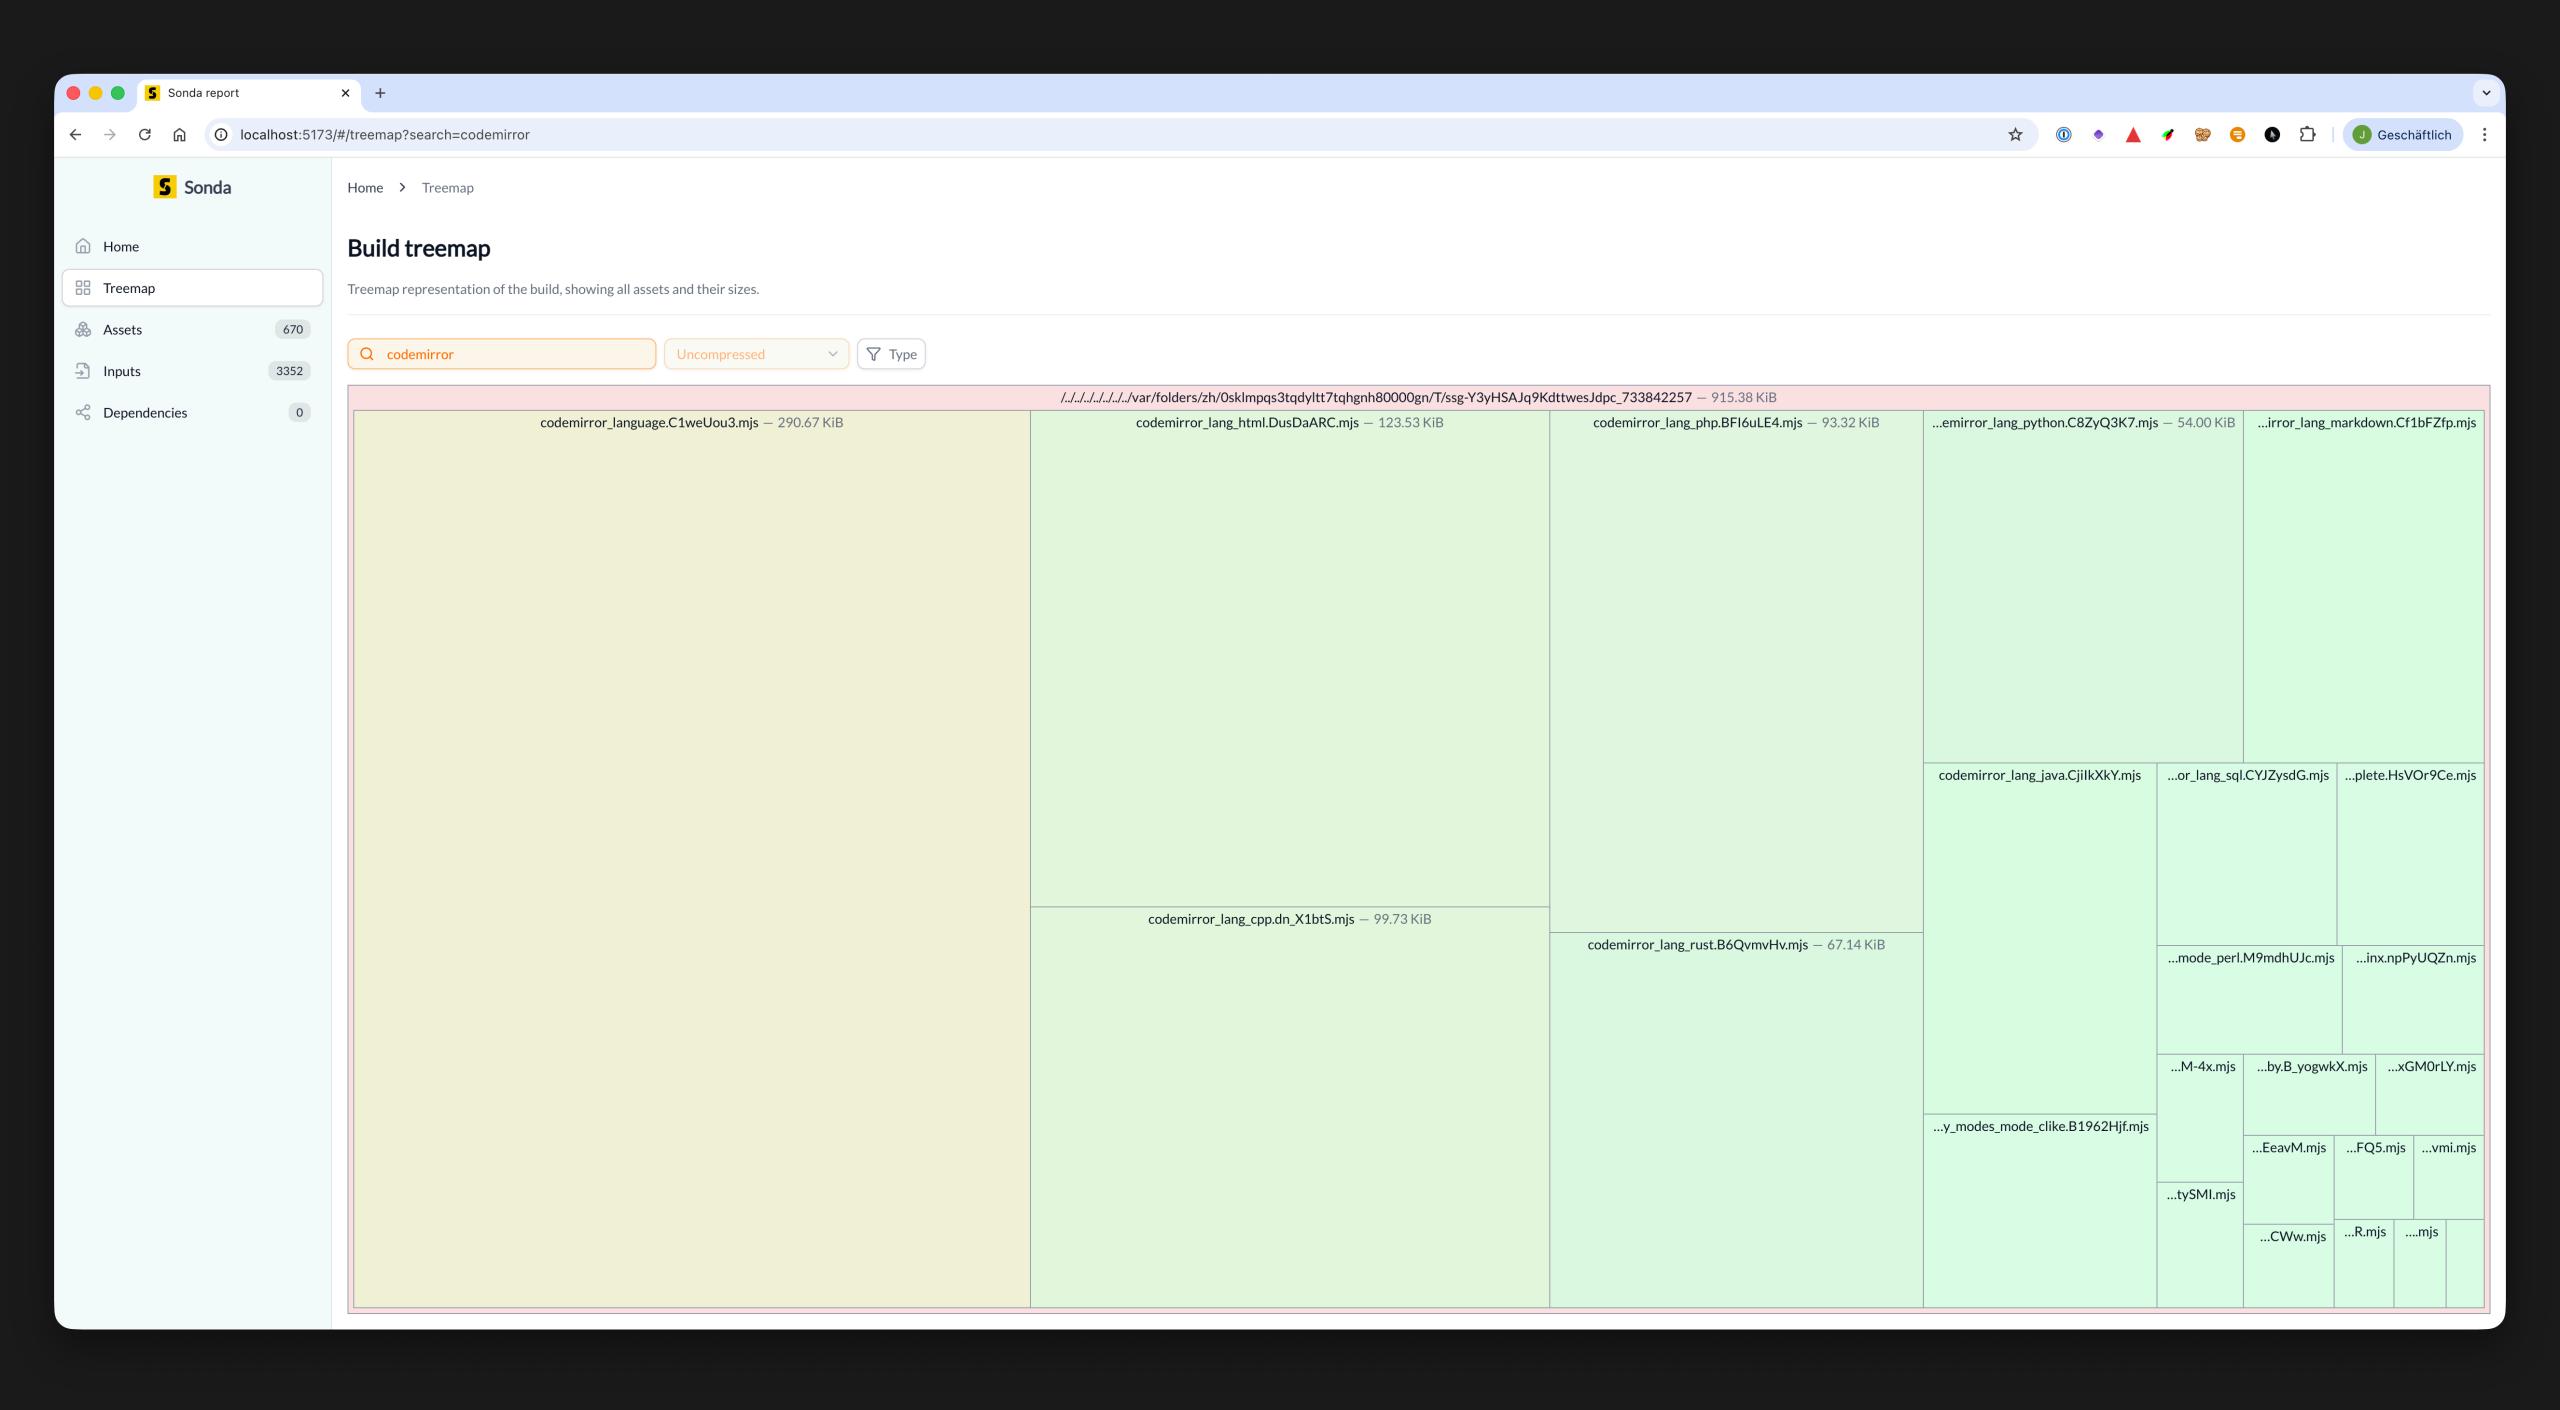
# Analyze your chunks

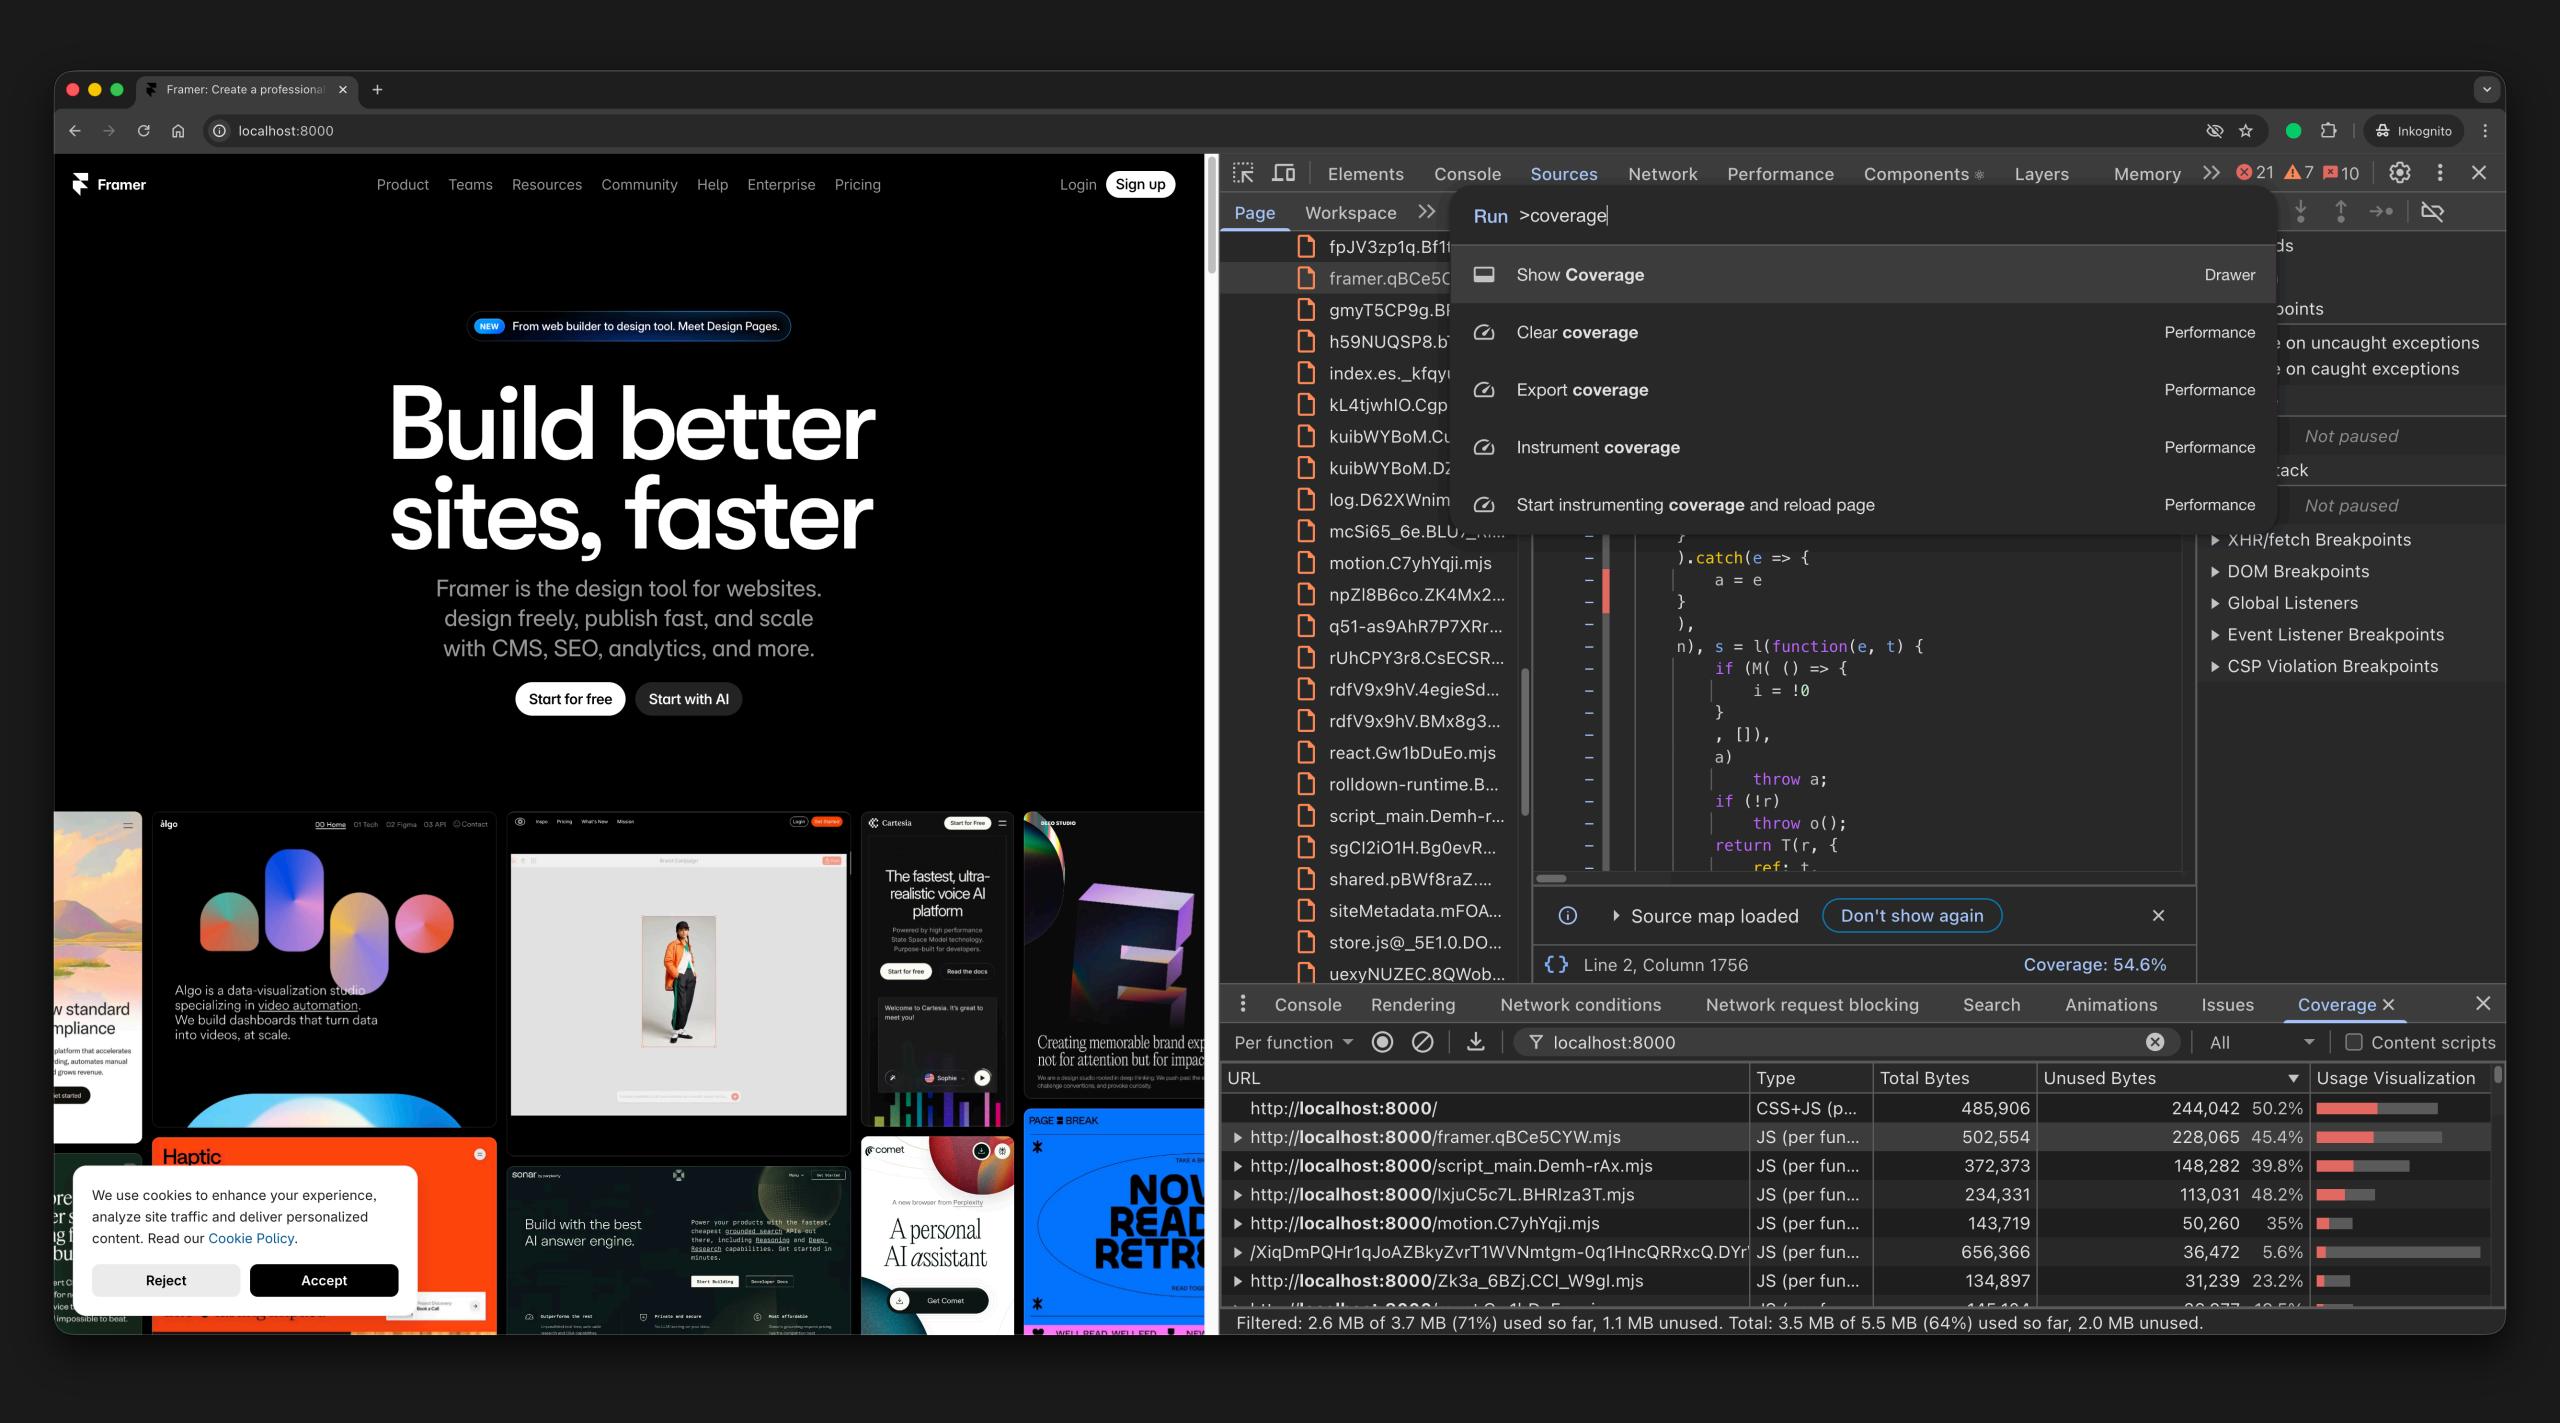
#### Analyzing bundles

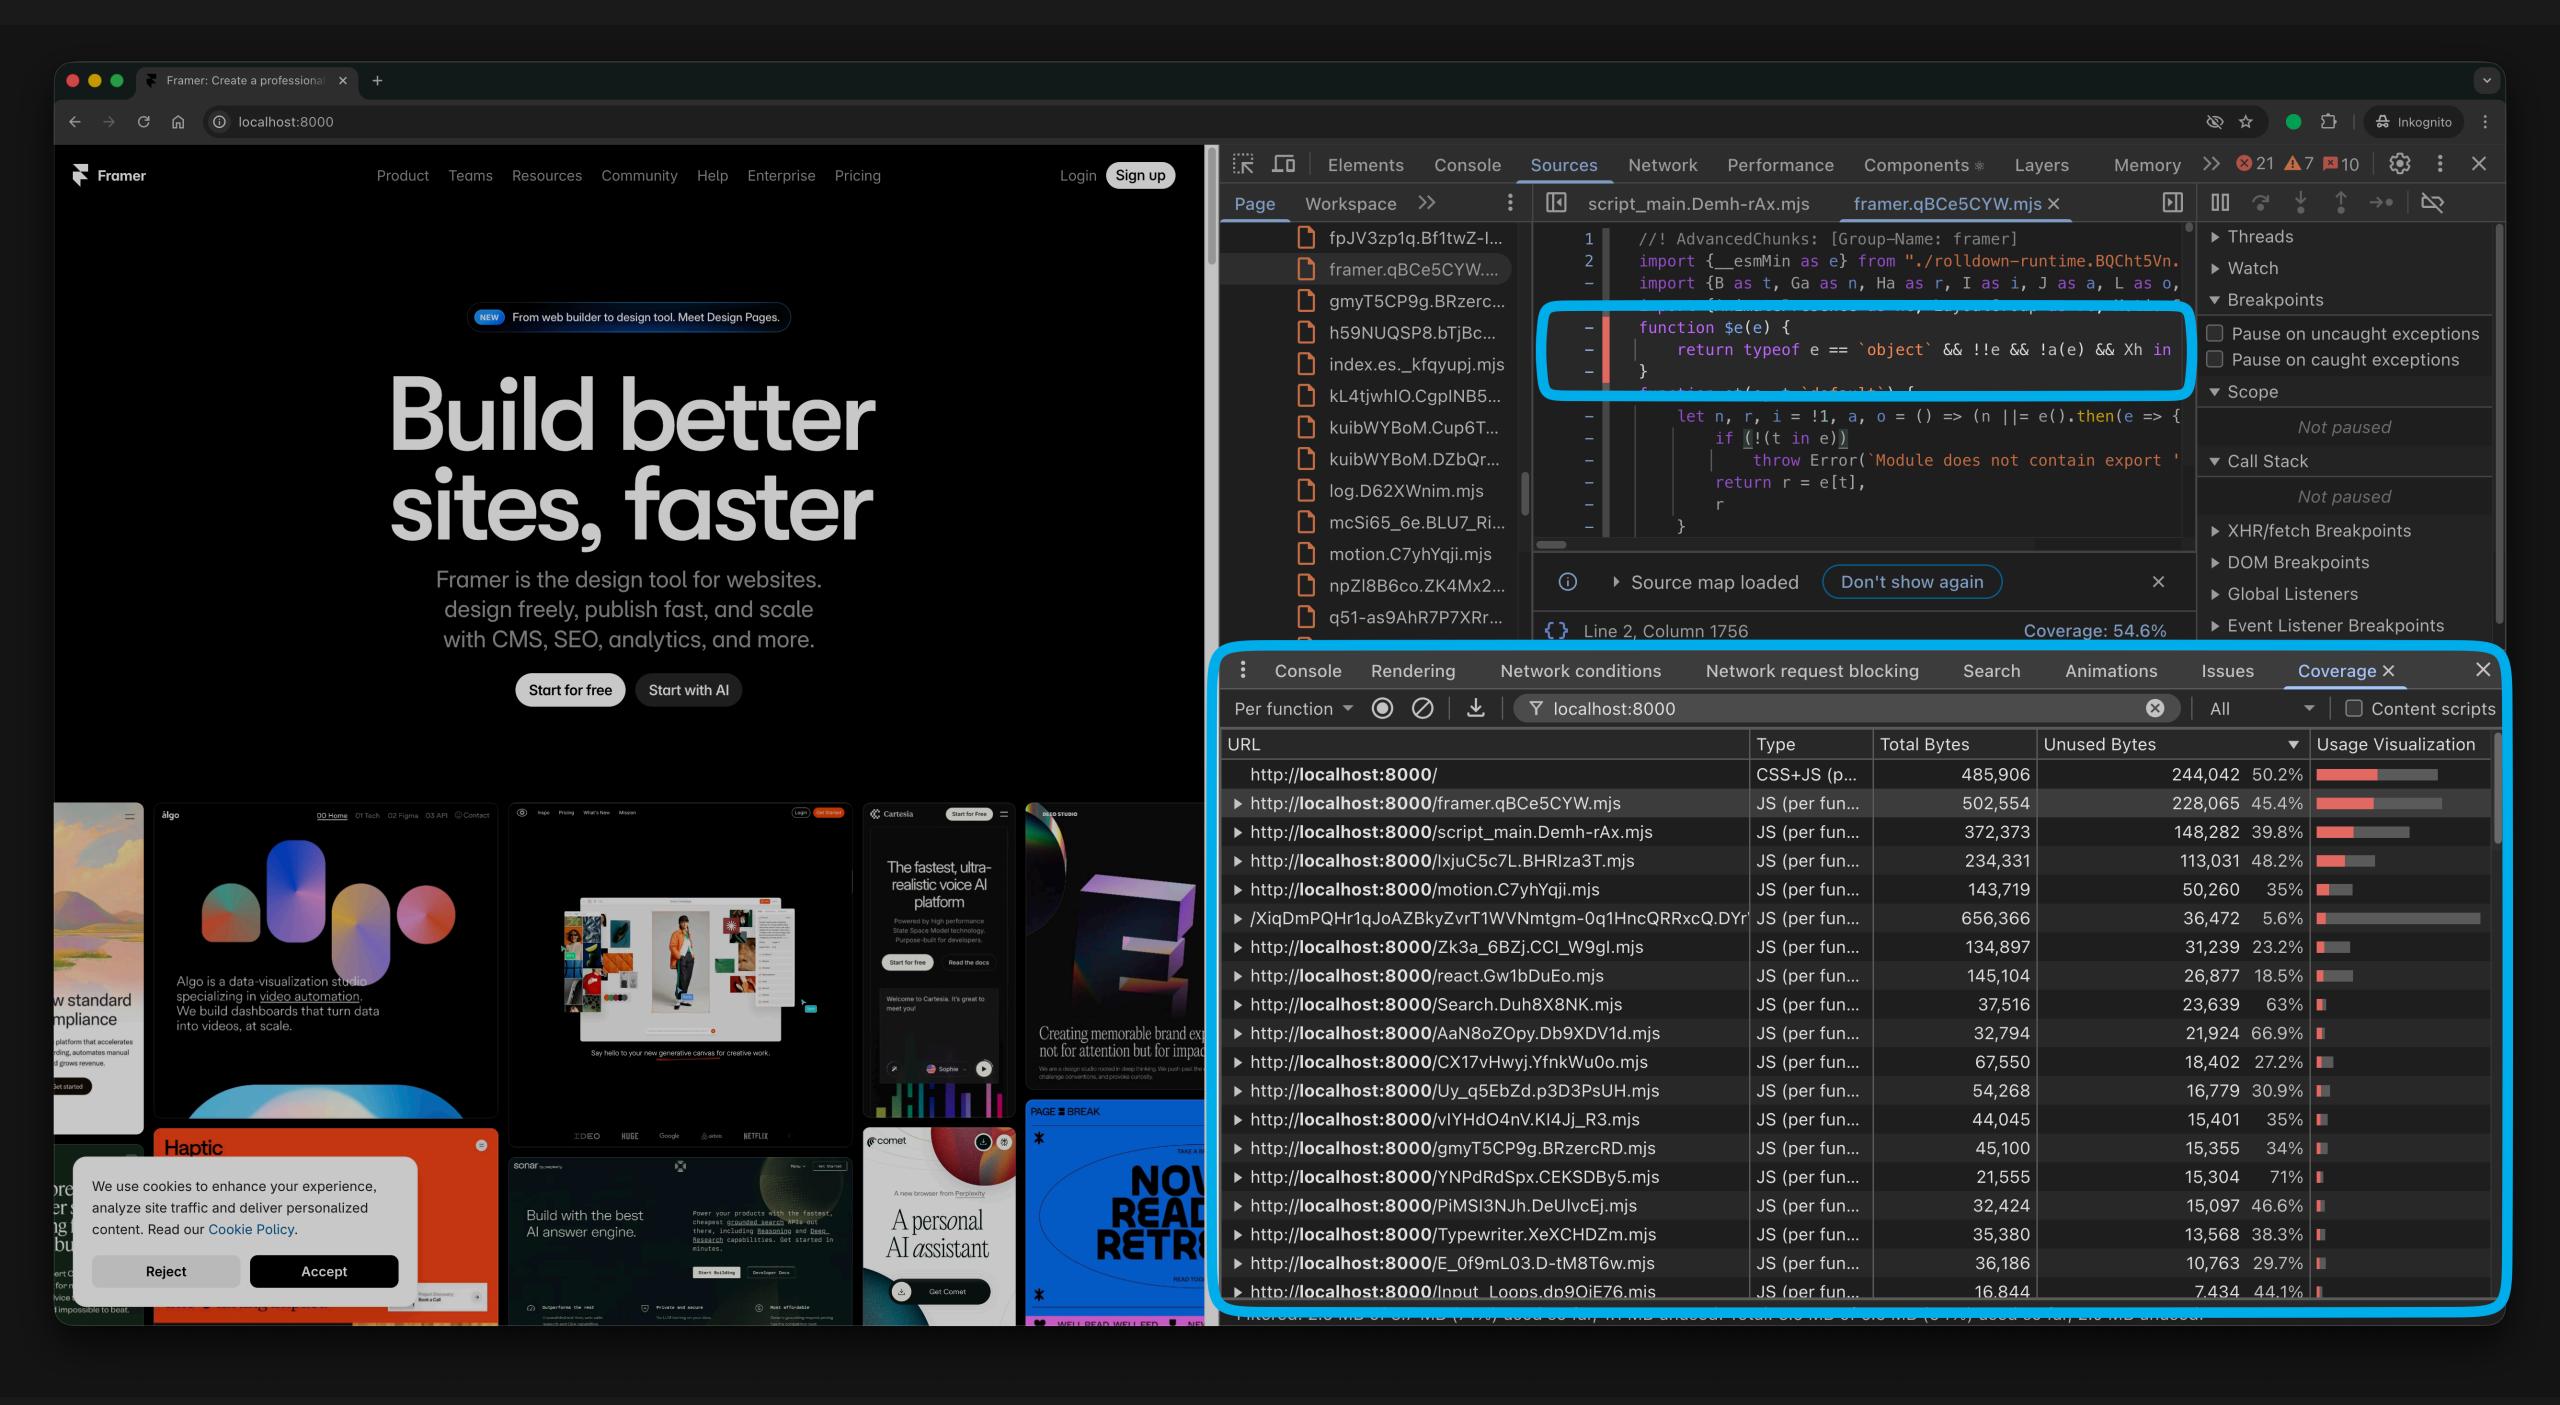
• Spoiler: Bundle analysis will be part of Vite Devtools, stay tuned for Anthony's talk

- For pure Rolldown users, currently there are two plugins:
  - Sonda
  - Rollup-plugin-visualizer
- DevTools is a viable option too









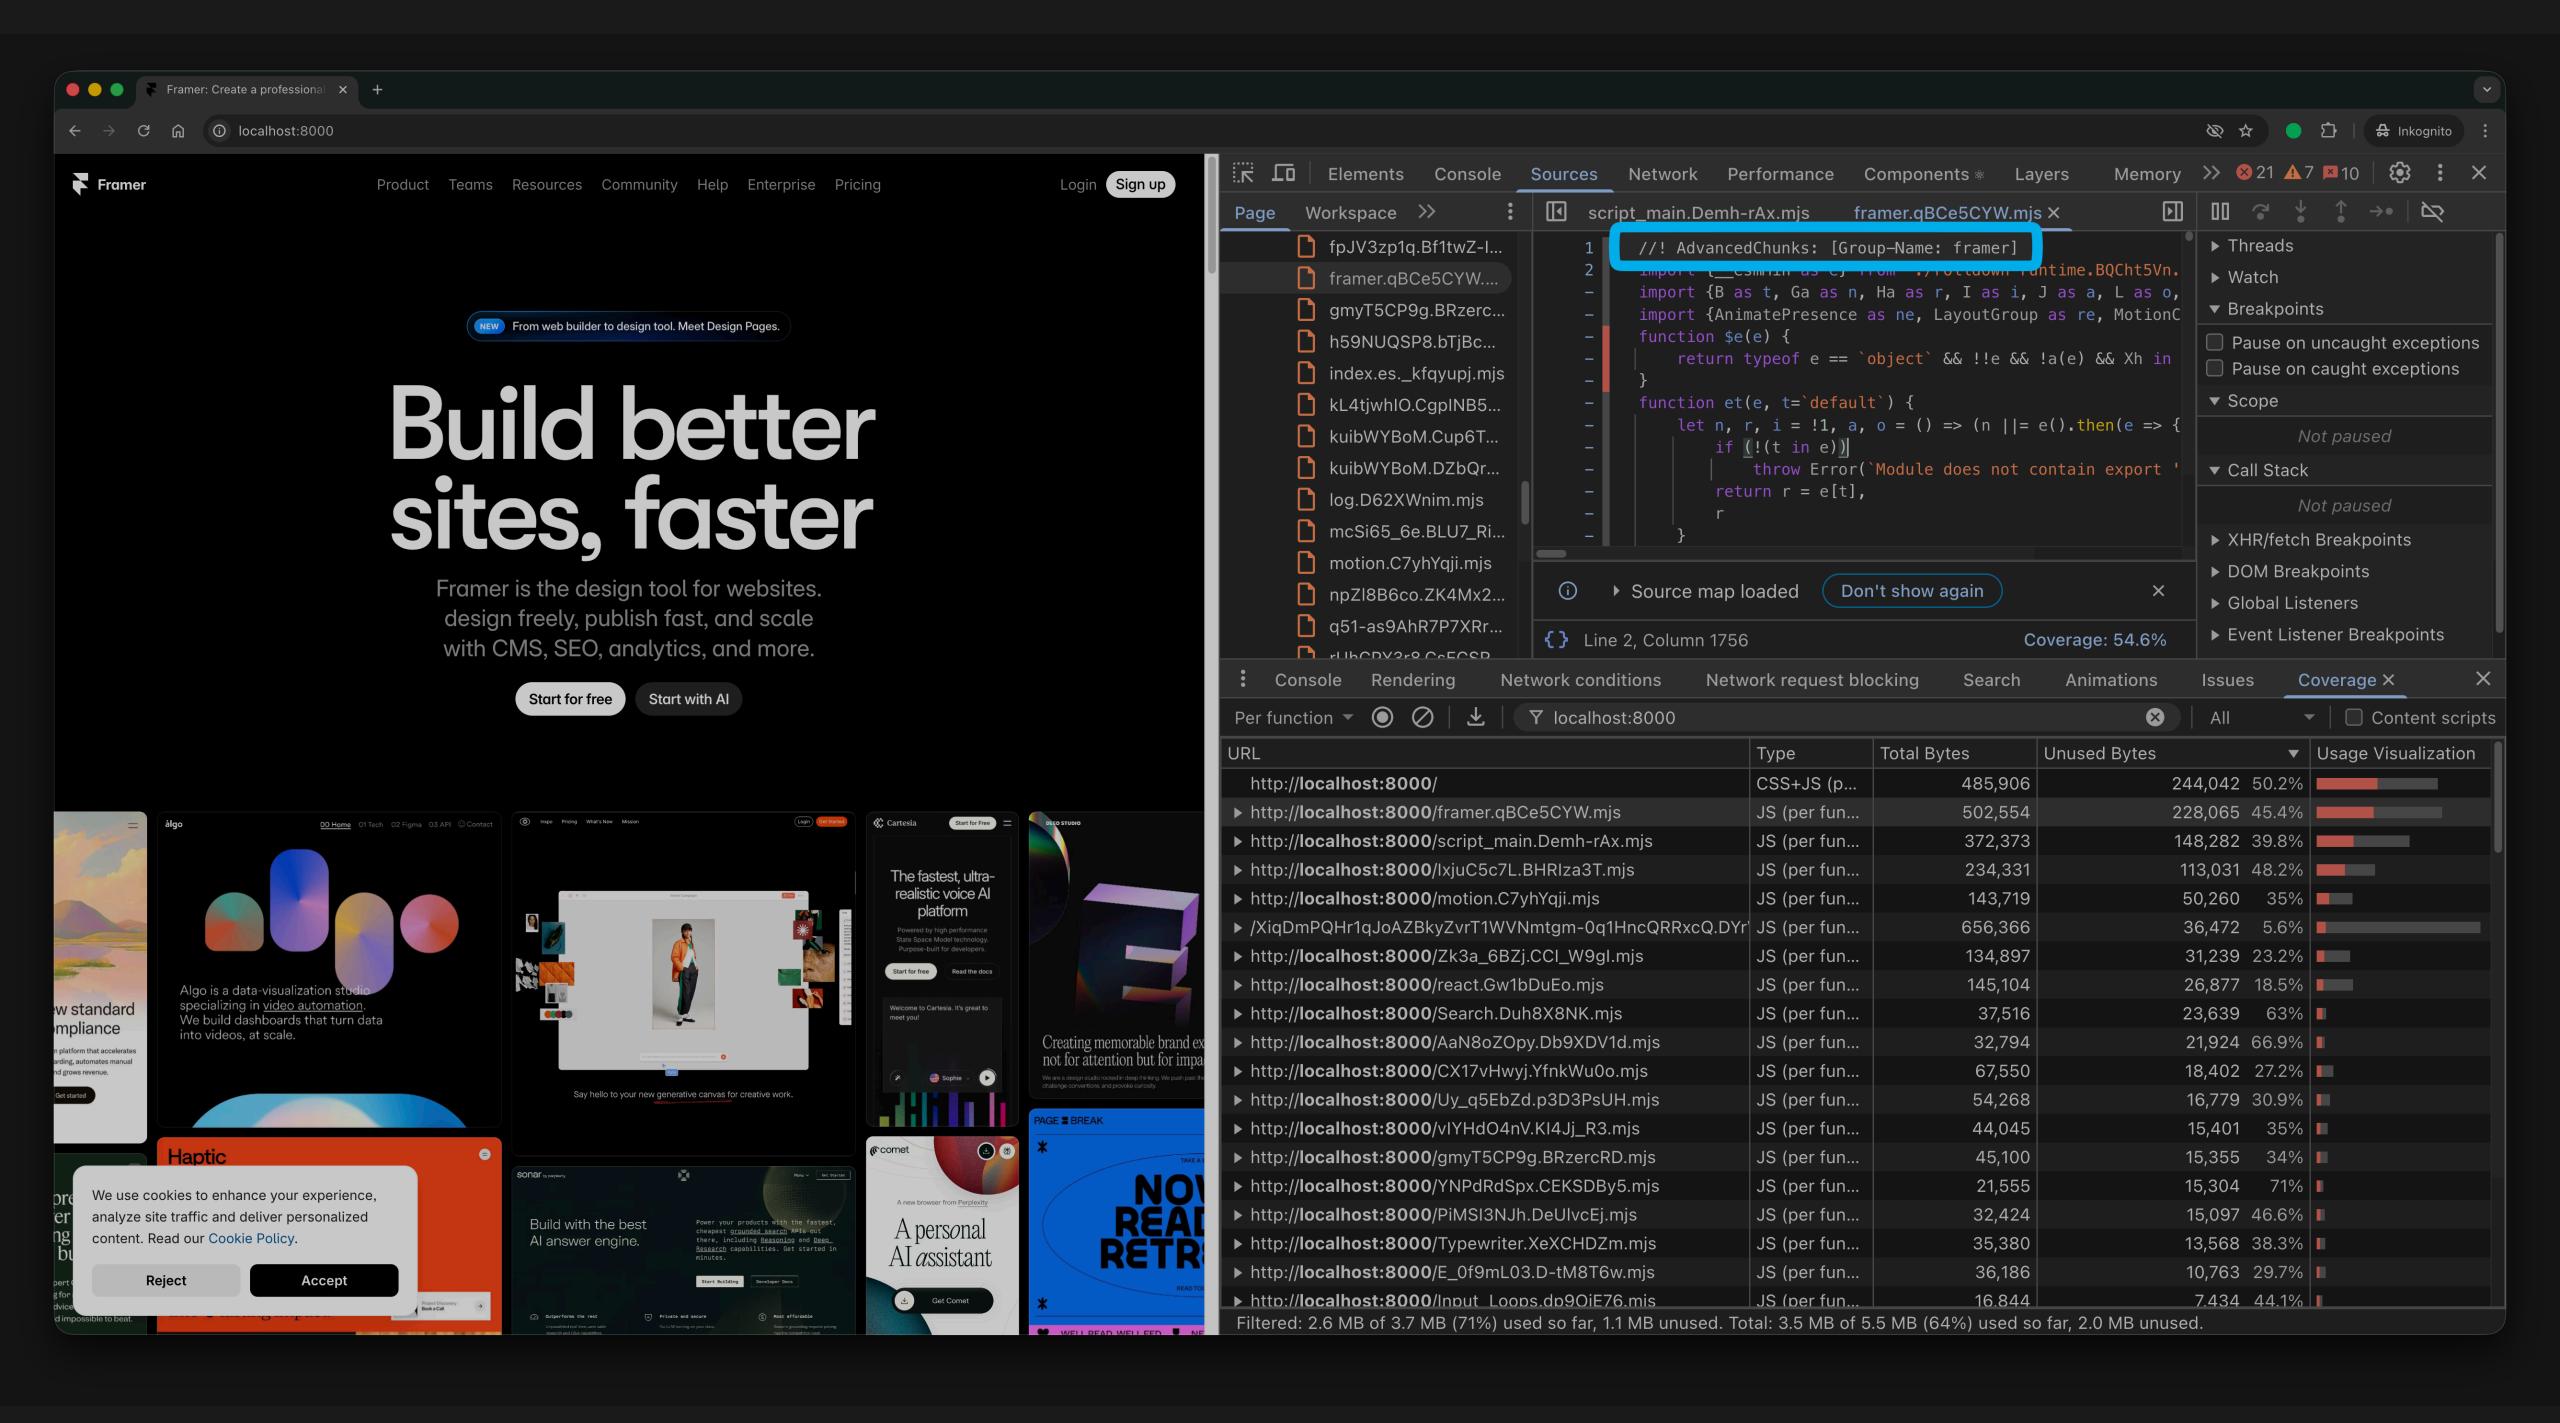
#### Analyzing bundles

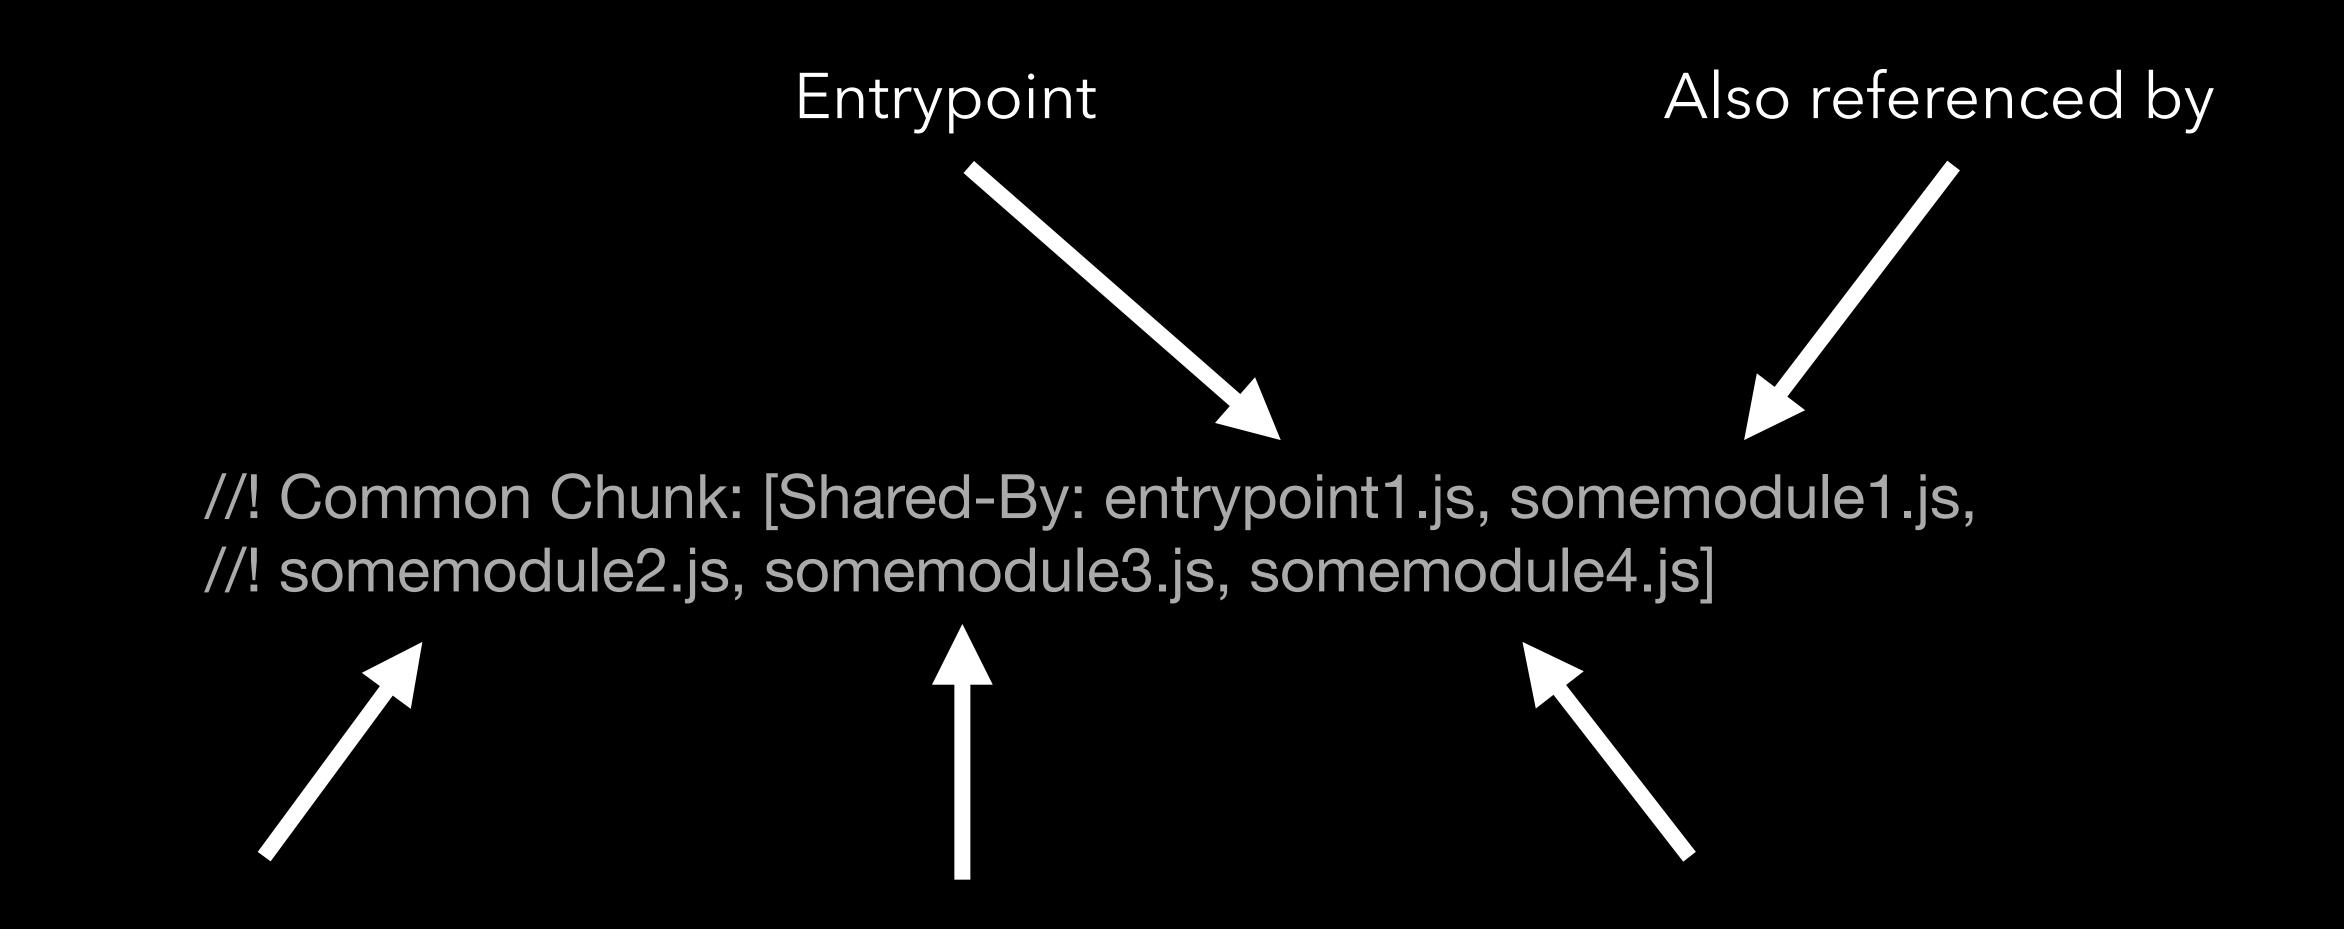
Follow discussions about emitting an esbuild-like metafile in #6425.

Once we have that, you can use:

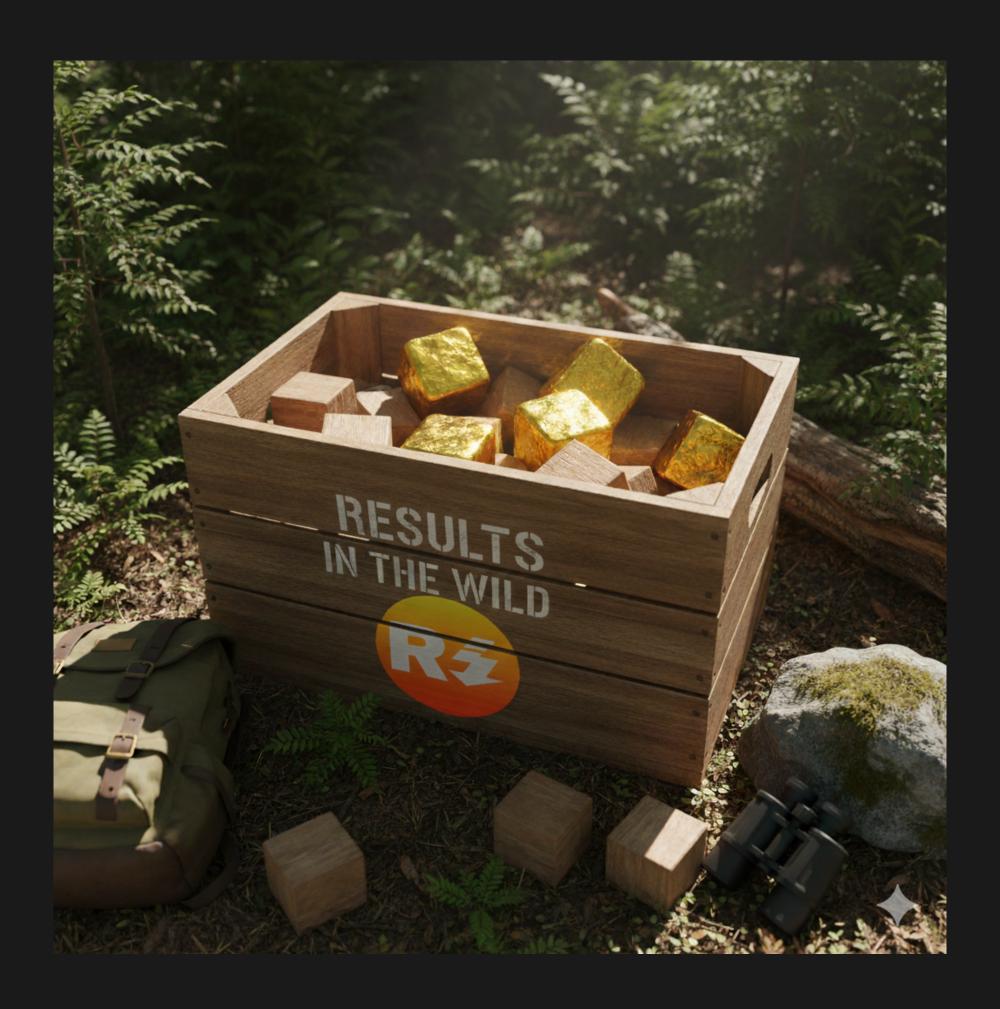
- hawkeye, bundle-buddy, esbuild-analyzer, Bundle Size Analyzer
- @rnx-kit/esbuild-bundle-analyzer
  - ...which unlocks support for <u>microsoft/webpack-bundle-compare</u> & 🌺 <u>Rsdoctor</u>

```
experimental: {
   attachDebugInfo: "full",
},
```





To get rid of this chunk  $\rightarrow$  make sure it's not imported by those modules



### Results

### Improvements to chunks

p25	20 → <b>14</b>
p75	67 → <b>22</b>
p90	80 → <b>30</b>
p99	95 → <b>54</b>

# 

Median drop in download size in both un- and compressed JS



Avg. improvement to LCP¹ for websites with 1-2MB of JS

# 

Avg. improvement to LCP¹ for websites with 2MB+ of JS.

## 

Faster LCP<sup>1</sup> at p90 across all Framer sites.

for slow devices & slow networks

#### A fast, stable and capable bundler is a win-win for everyone.

For us, our customers & their visitors, the Vite community, the web.



### Thank you!

Questions?

@kurtextrem on X, LinkedIn, BlueSky

https://kurtextrem.de