

Render Performance in the Age of AI

A Lesson in Controlling Entropy

Jenna Zeigen

NY UX Speed and Web Performance Meetup

June 18, 2026

hi, i'm jenna

- Former Senior Staff perf girlie @ Notion & Slack
- Speaker Manager for this meetup!
- Find me on the internet at @zeigenvector
- Find these slides at jenna.is/at-ny-perf-2026

hire me??



how did **AI tools** change how

we kept "Annotation"'s

render performance and

interaction latency in check?

why "do" performance?

- ~~So the graphs go in the right direction?~~
 - ~~So we make more money??~~
 - ~~So people don't write bad things about us on Hacker News???~~
- ✨ So our users have a ~~great pleasant~~
not-bad experience! ✨

how does AI change things?

- AI amplifies patterns, both good and bad
- the patterns are from both your codebase and the *internet*
- we're a lot more "productive" now, and more code means more opportunities for mistakes

how does AI change things?

- lower barrier to entry for contributing to the codebase means people don't always understand the implications of the changes they are making

how does AI change things?

Performance is about tradeoffs,
but AI cannot be taught "taste" or
"it depends" effectively

how *doesn't* AI change things?

- performance problems ultimately come down to the computer doing more work than it has to
- often, the result of systemic issues
- to solve this, "you" have to understand the system
- "you" still need to understand what's currently slow

step 0: measure!

you still need metrics

- What metrics make the most sense for **your** application?
- Core Web Vitals could still be it
- Does AI change the way your users will be interacting with the product, and therefore the ways it could be slow?

you still need metrics

- INP = clicking + typing



jenna Oct 30th, 2025 at 3:33 PM

Not really an investigation, but... has there been any recommendations about INP for typing heavy applications? 200ms INP keystrokes seems like it's going to be a bad time for users, but canonical INP lore says <200ms is good. For context, our "clicky" INP p75 is 175ms but our keypress INP's is 40ms. Even curious what people's vibes on this are! (edited)



Barry Pollard Oct 30th, 2025 at 4:29 PM

200ms INP keystrokes seems like it's going to be a bad time for users, but canonical INP lore says <200ms is good. Core Web Vitals targets are broad, high-level, and aimed to apply to all pages. They were also set with one eye to achievability: <https://web.dev/articles/defining-core-web-vitals-thresholds>.

I agree for typing, 200ms is often way too high and you should aim for lower. But, for simplicity we don't have different thresholds for typing and button clicks—which are two very different interactions.

Exactly what the target time should be, I don't know.


Less than 100ms is generally considered imperceptible to most people so maybe that's a good target? Though there's other research saying that's not a fixed threshold and people are different.

You could consider typing skills. Ideally you'd show each letter as it's typed. A good Words Per Minute (WPM) is 60, and advanced is 100. Let's take the 100 example cause even non-advanced typists can reach high speed for certain typing, even if they can't sustain it like an advanced typist could. A "word" is generally considered to be 5 letters for typing measurement so that means we need $100/60 = 1.667$ words per second or 8.3333 letters per second. Which means you need to respond to a letter being typed every $1000/8.333 = 120$ milliseconds, if I've got my calculations correct.

So for two different reasons, I'd aim for 100-120ms. But that's a theoretical limit. It would be best tested with some people (especially typists) to see if they feel it's responsive. (edited)



it's a metric!



"~23% of users are experiencing
a >500ms interaction"

my talk two years ago

- 1. Targeting Problem Components**
- 2. Broad-Spectrum Solutions**



my talk today!

- 1. Targeting Problem Components**
- 2. Broad-Spectrum Solutions**

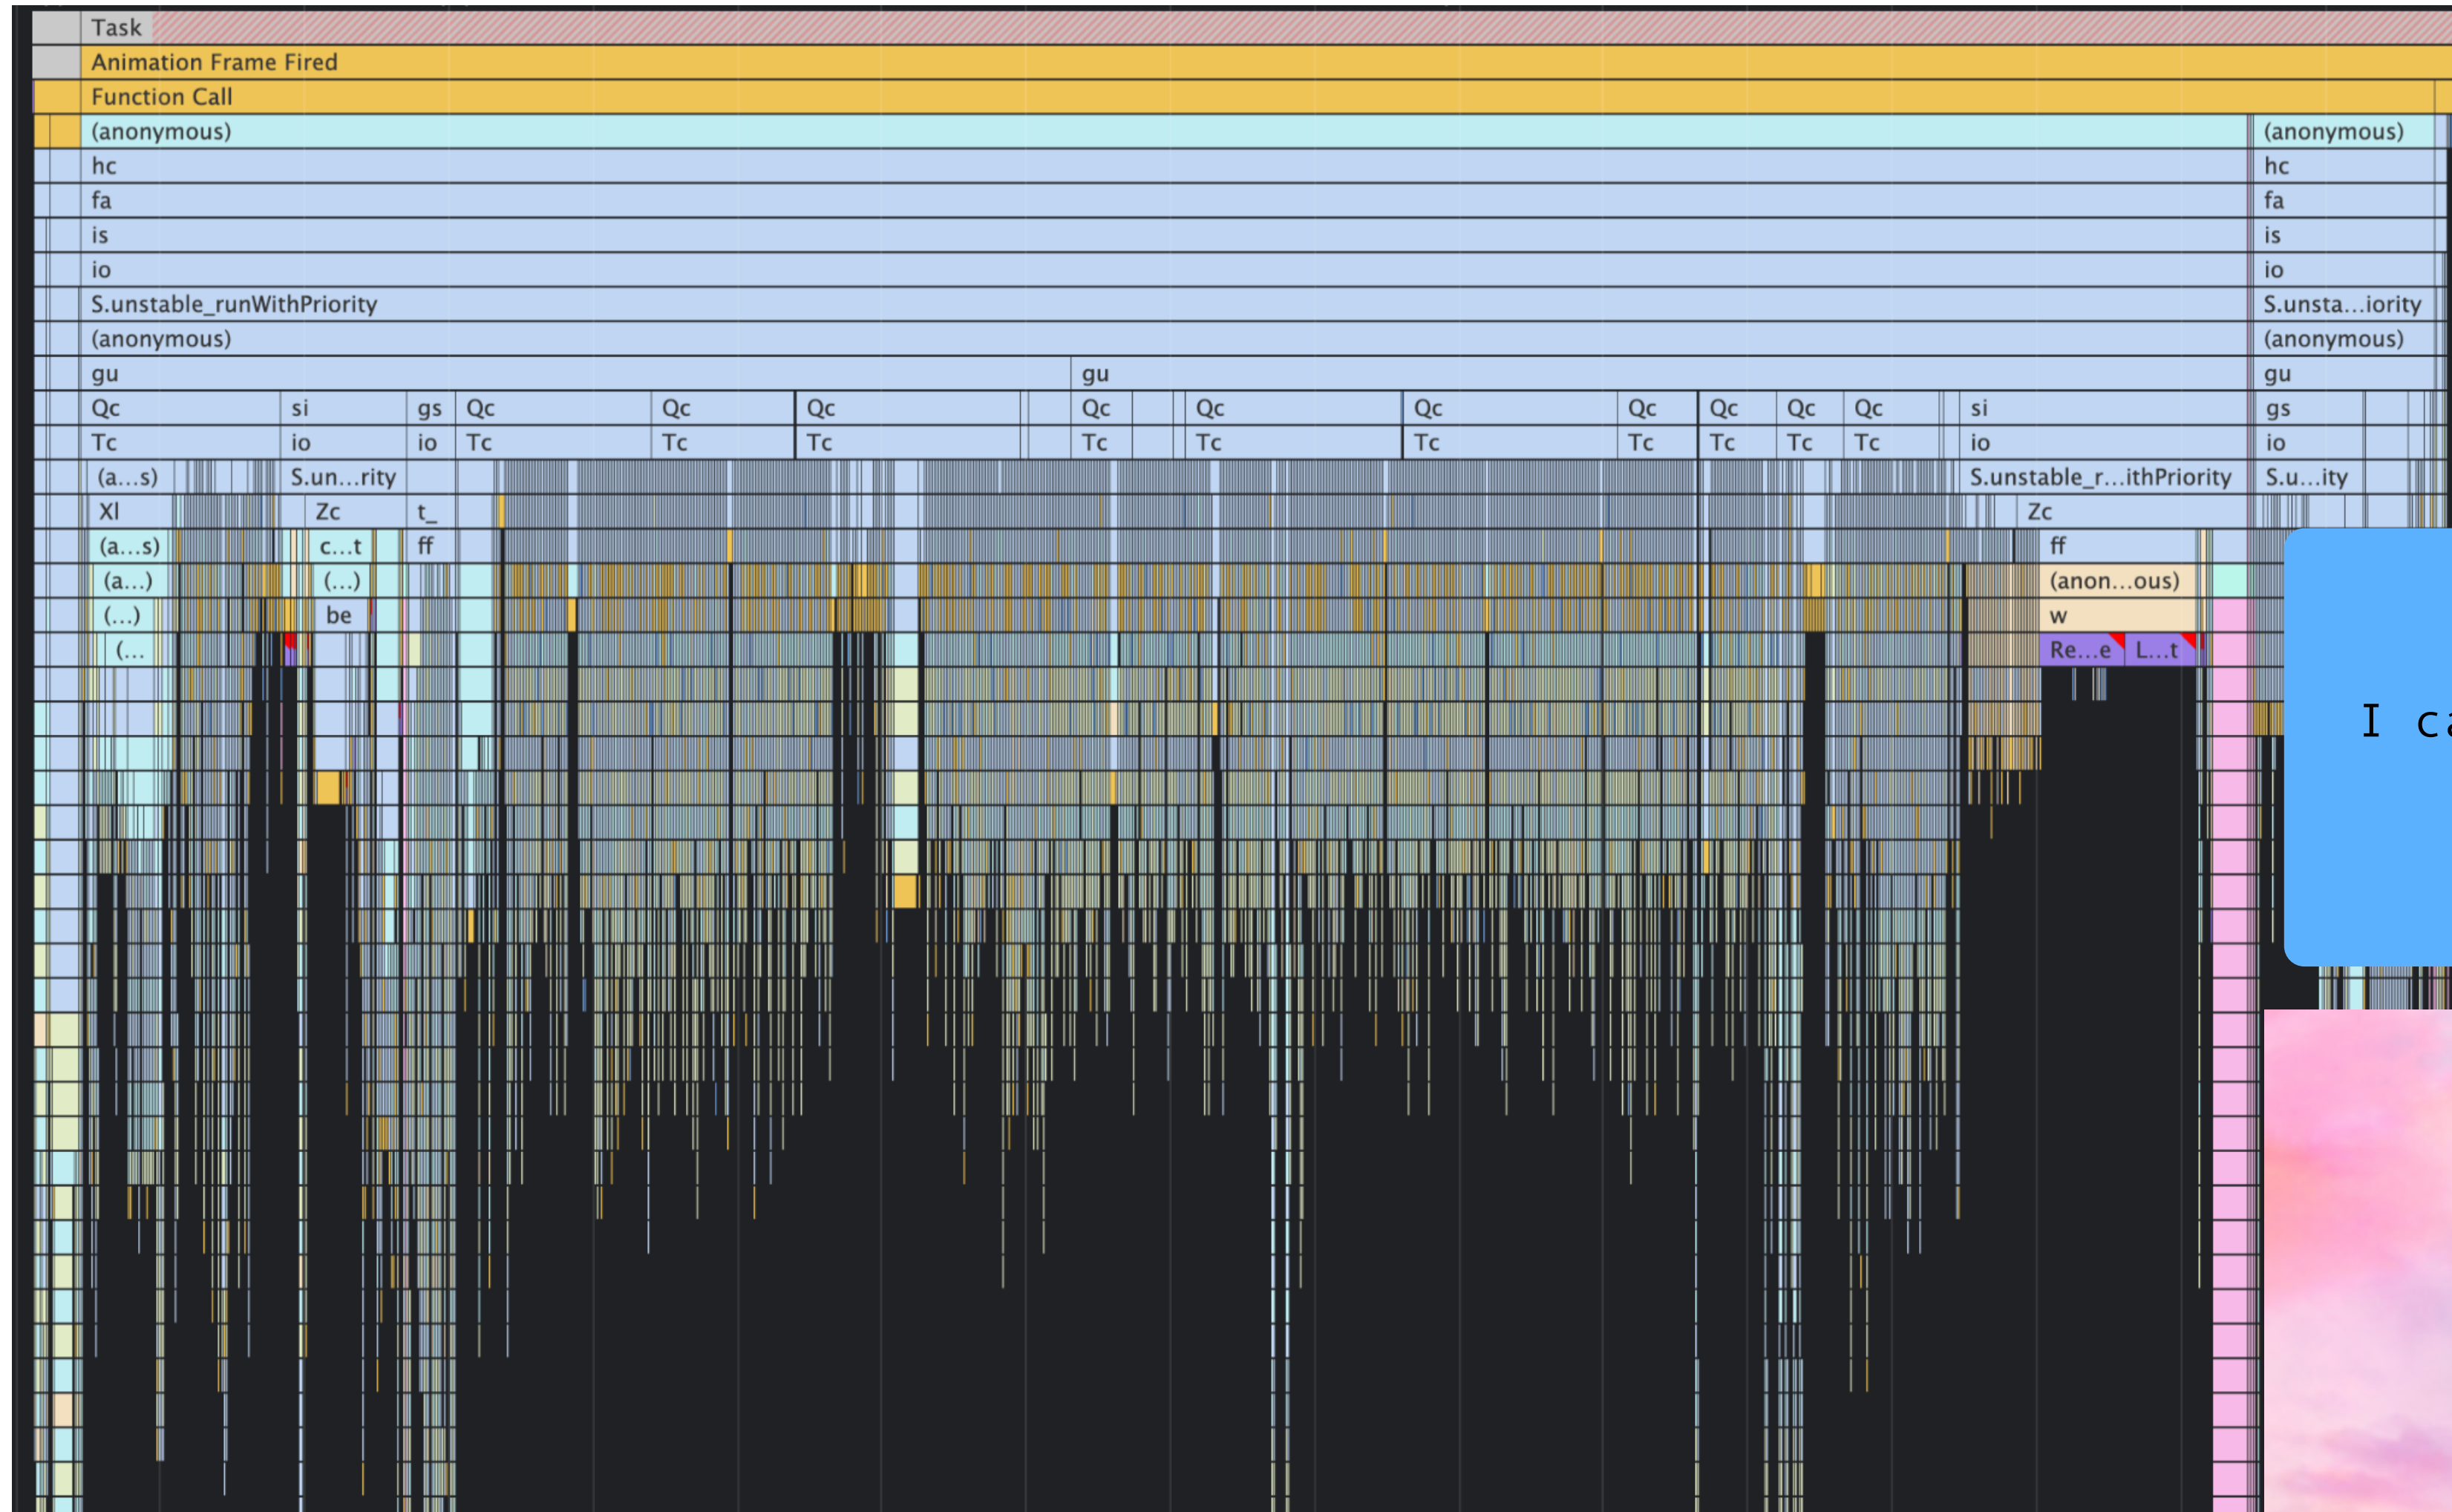
3. Diagnosing Regressions



broad spectrum

solutions

diagnosis: ~~murder~~ render



I can't pretend it's okay when it's not
It's death by a thousand cuts

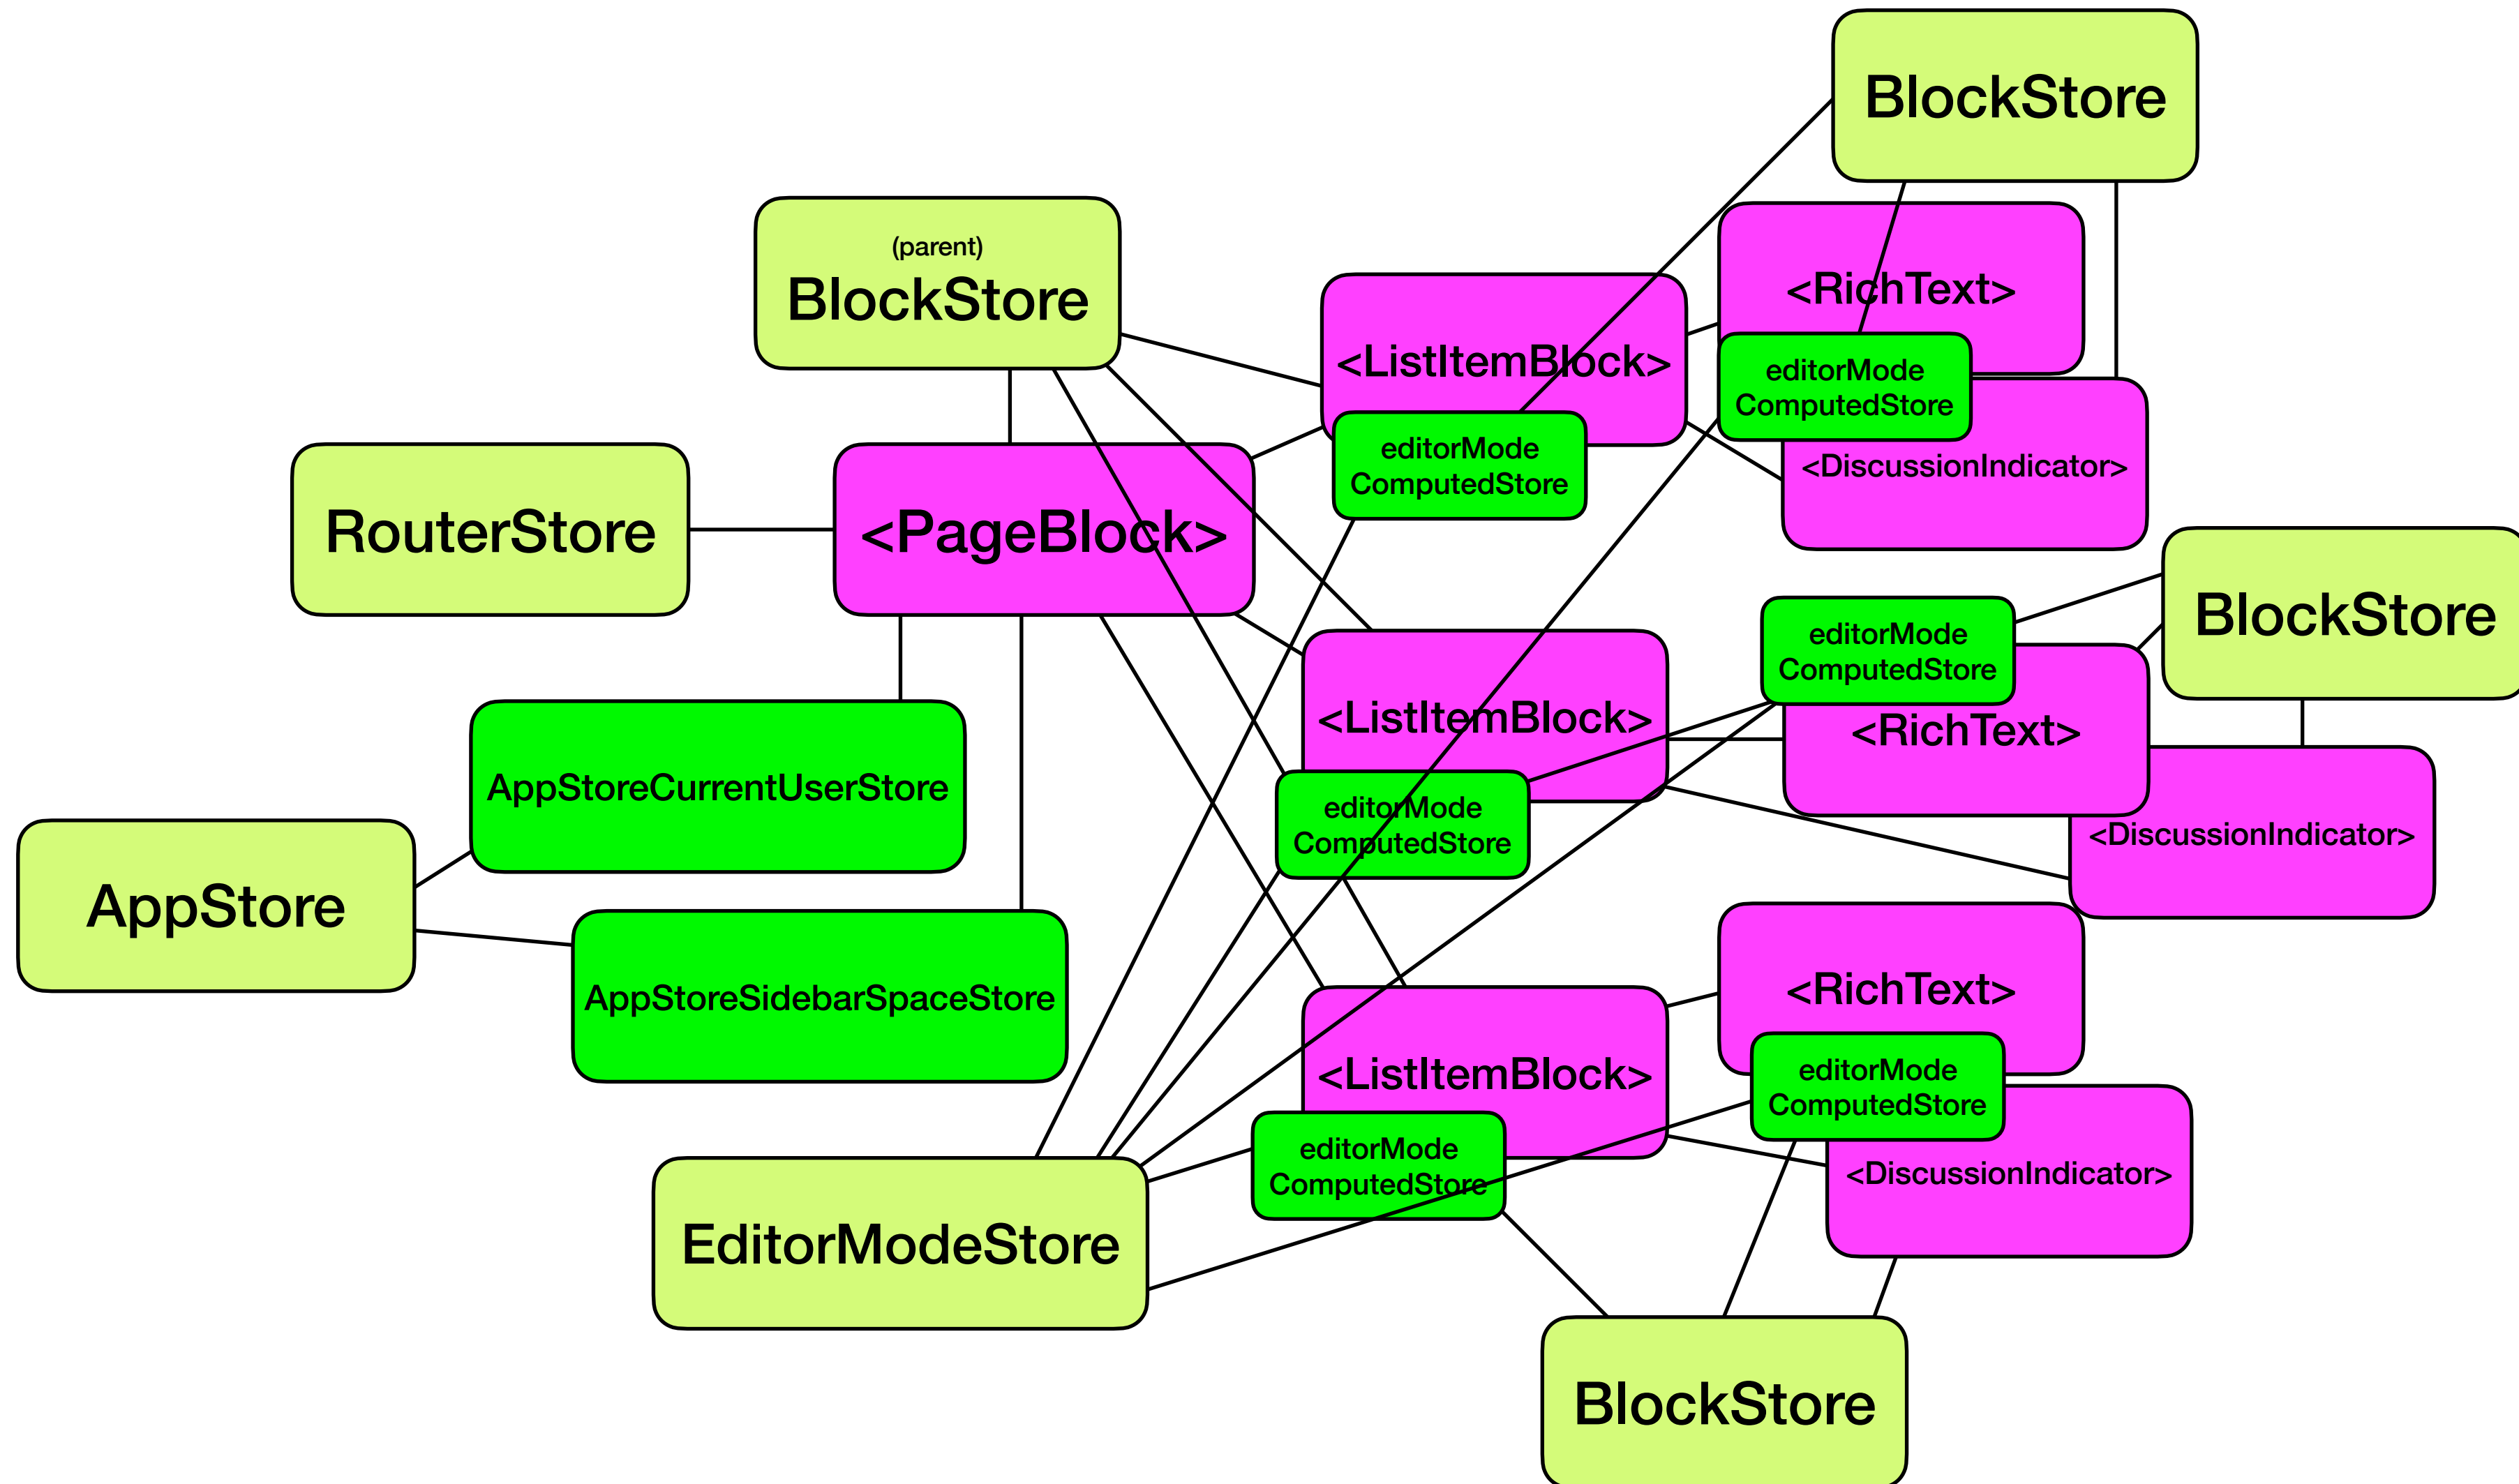


p.s. this is actually a profile of Slack from like 5 years ago
but the papercuttiness is the same!!!!

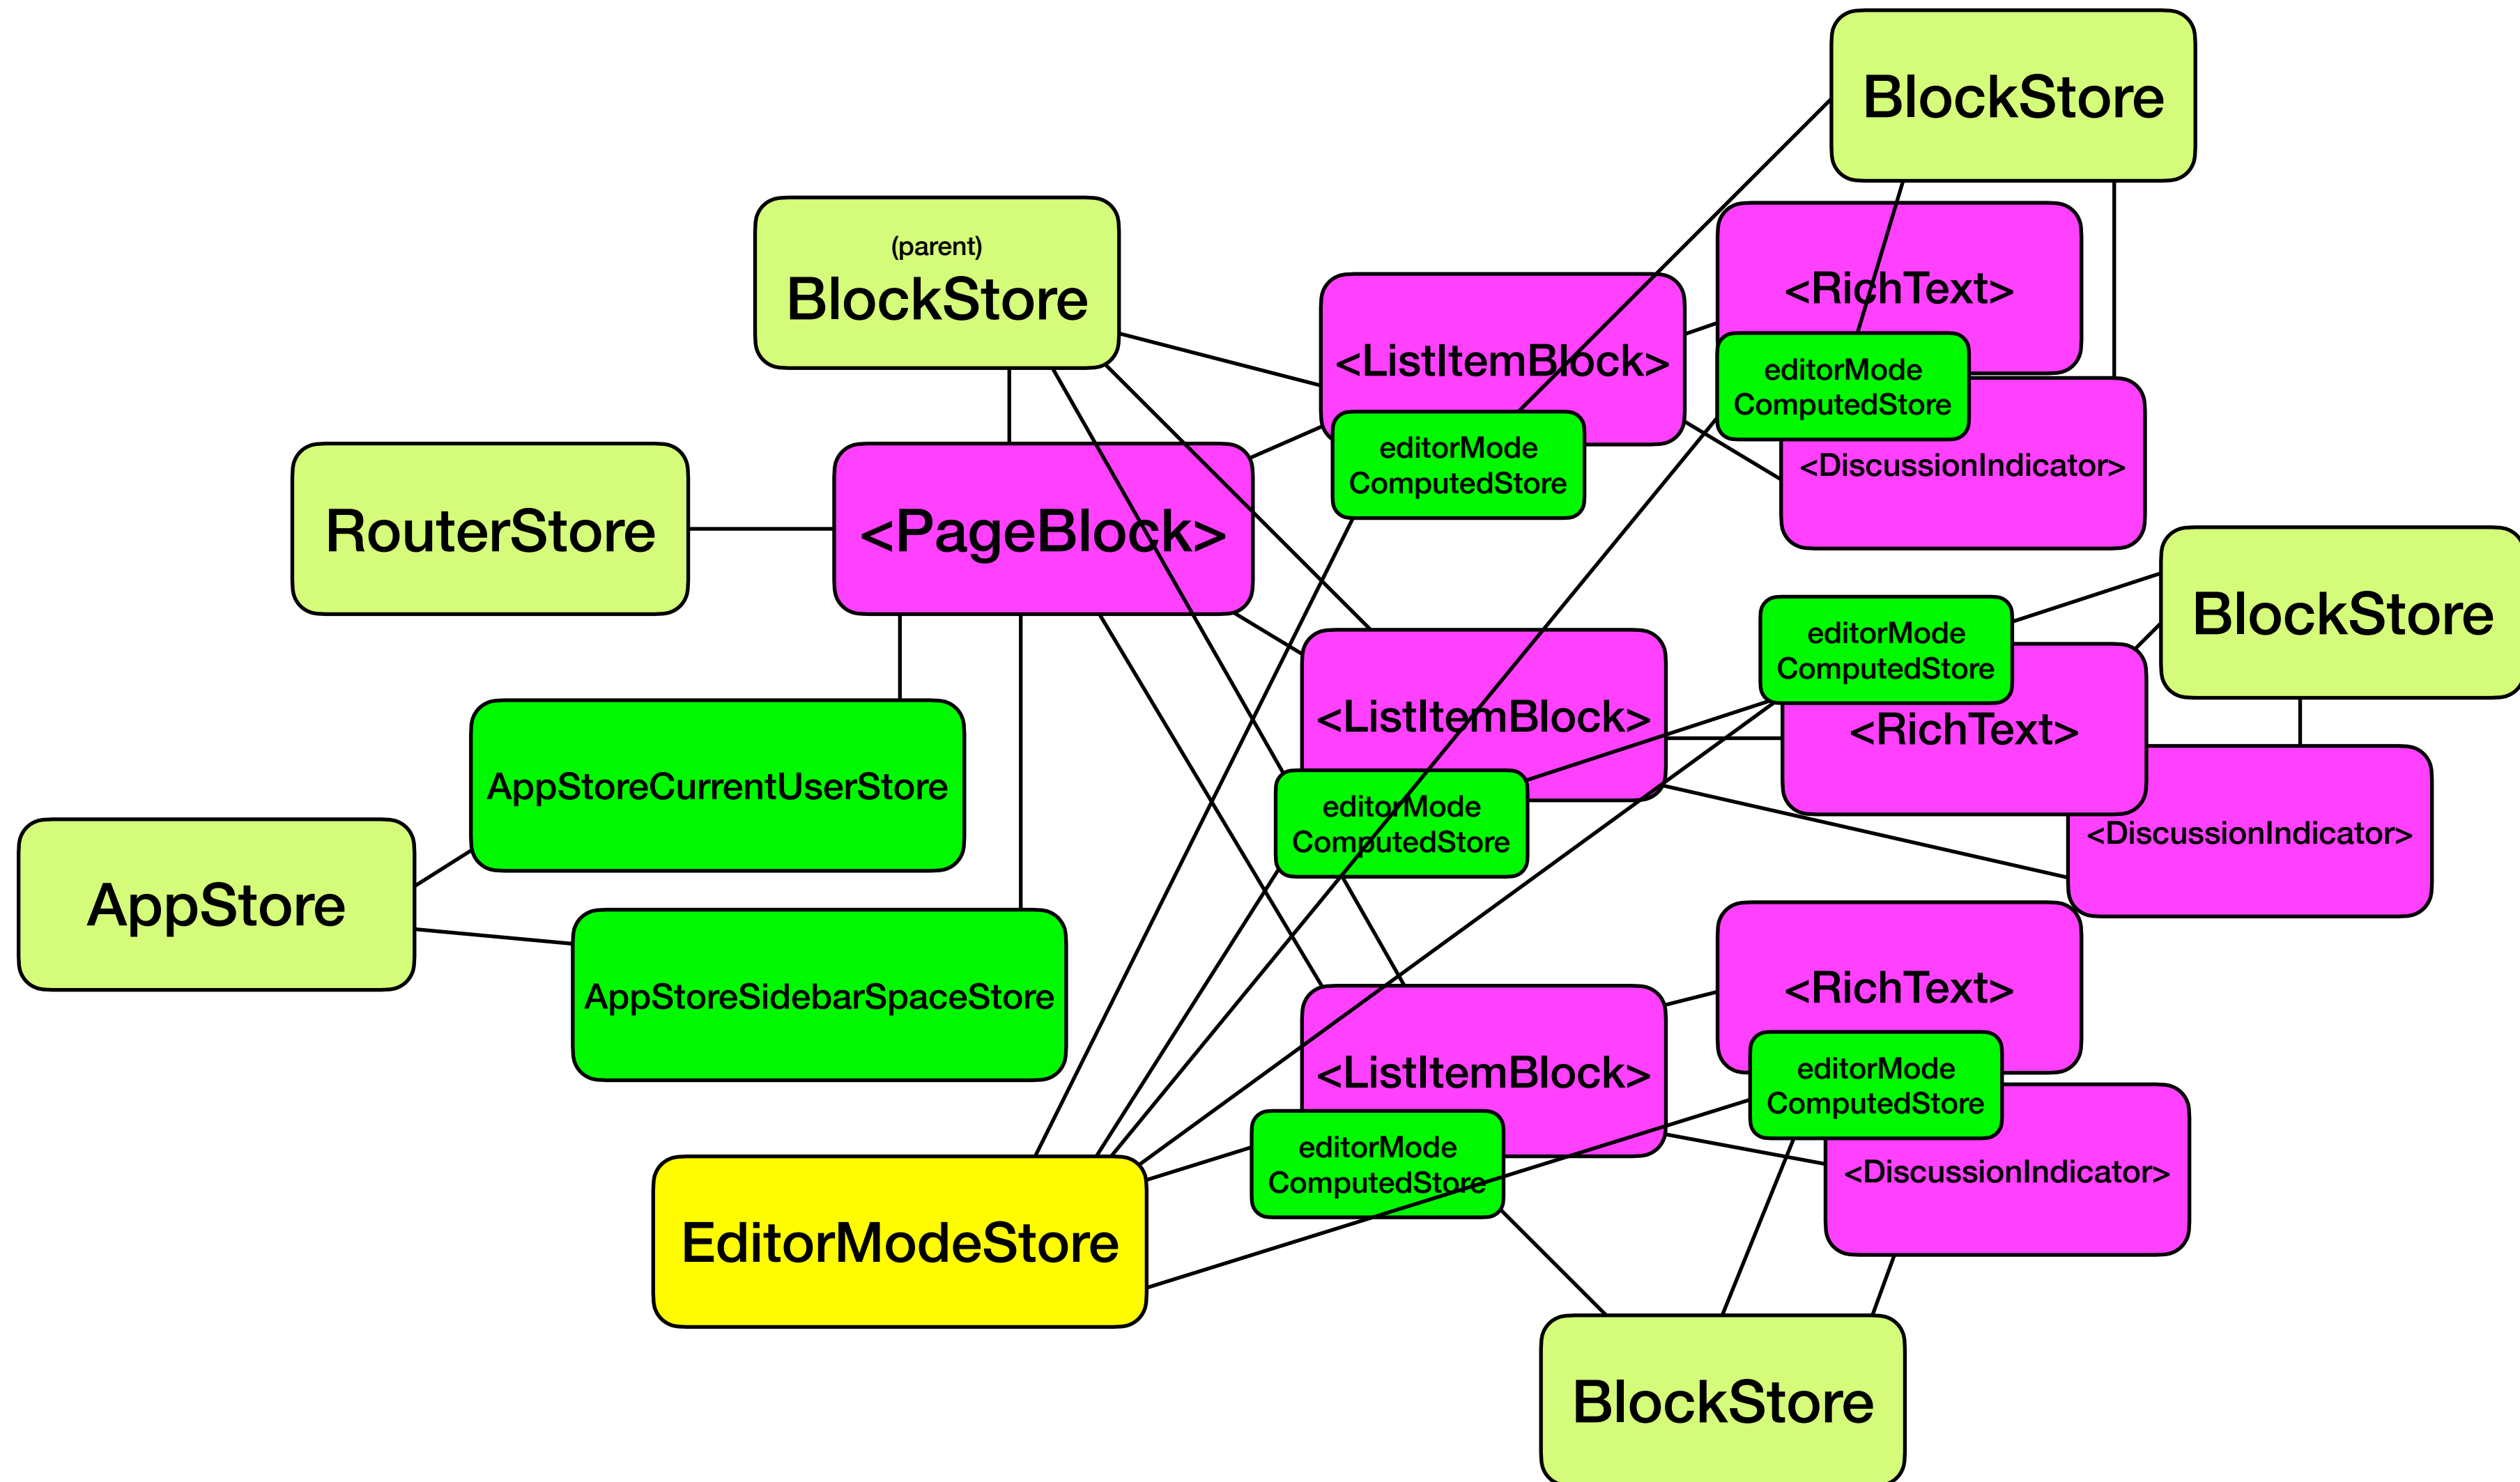
diagnosis: ~~murder~~ render

- how do the frameworks or systems your app is built on work?
- can the AI help you better understand it?
- can you help the AI better understand it?
- documentation can be read by humans and robots!

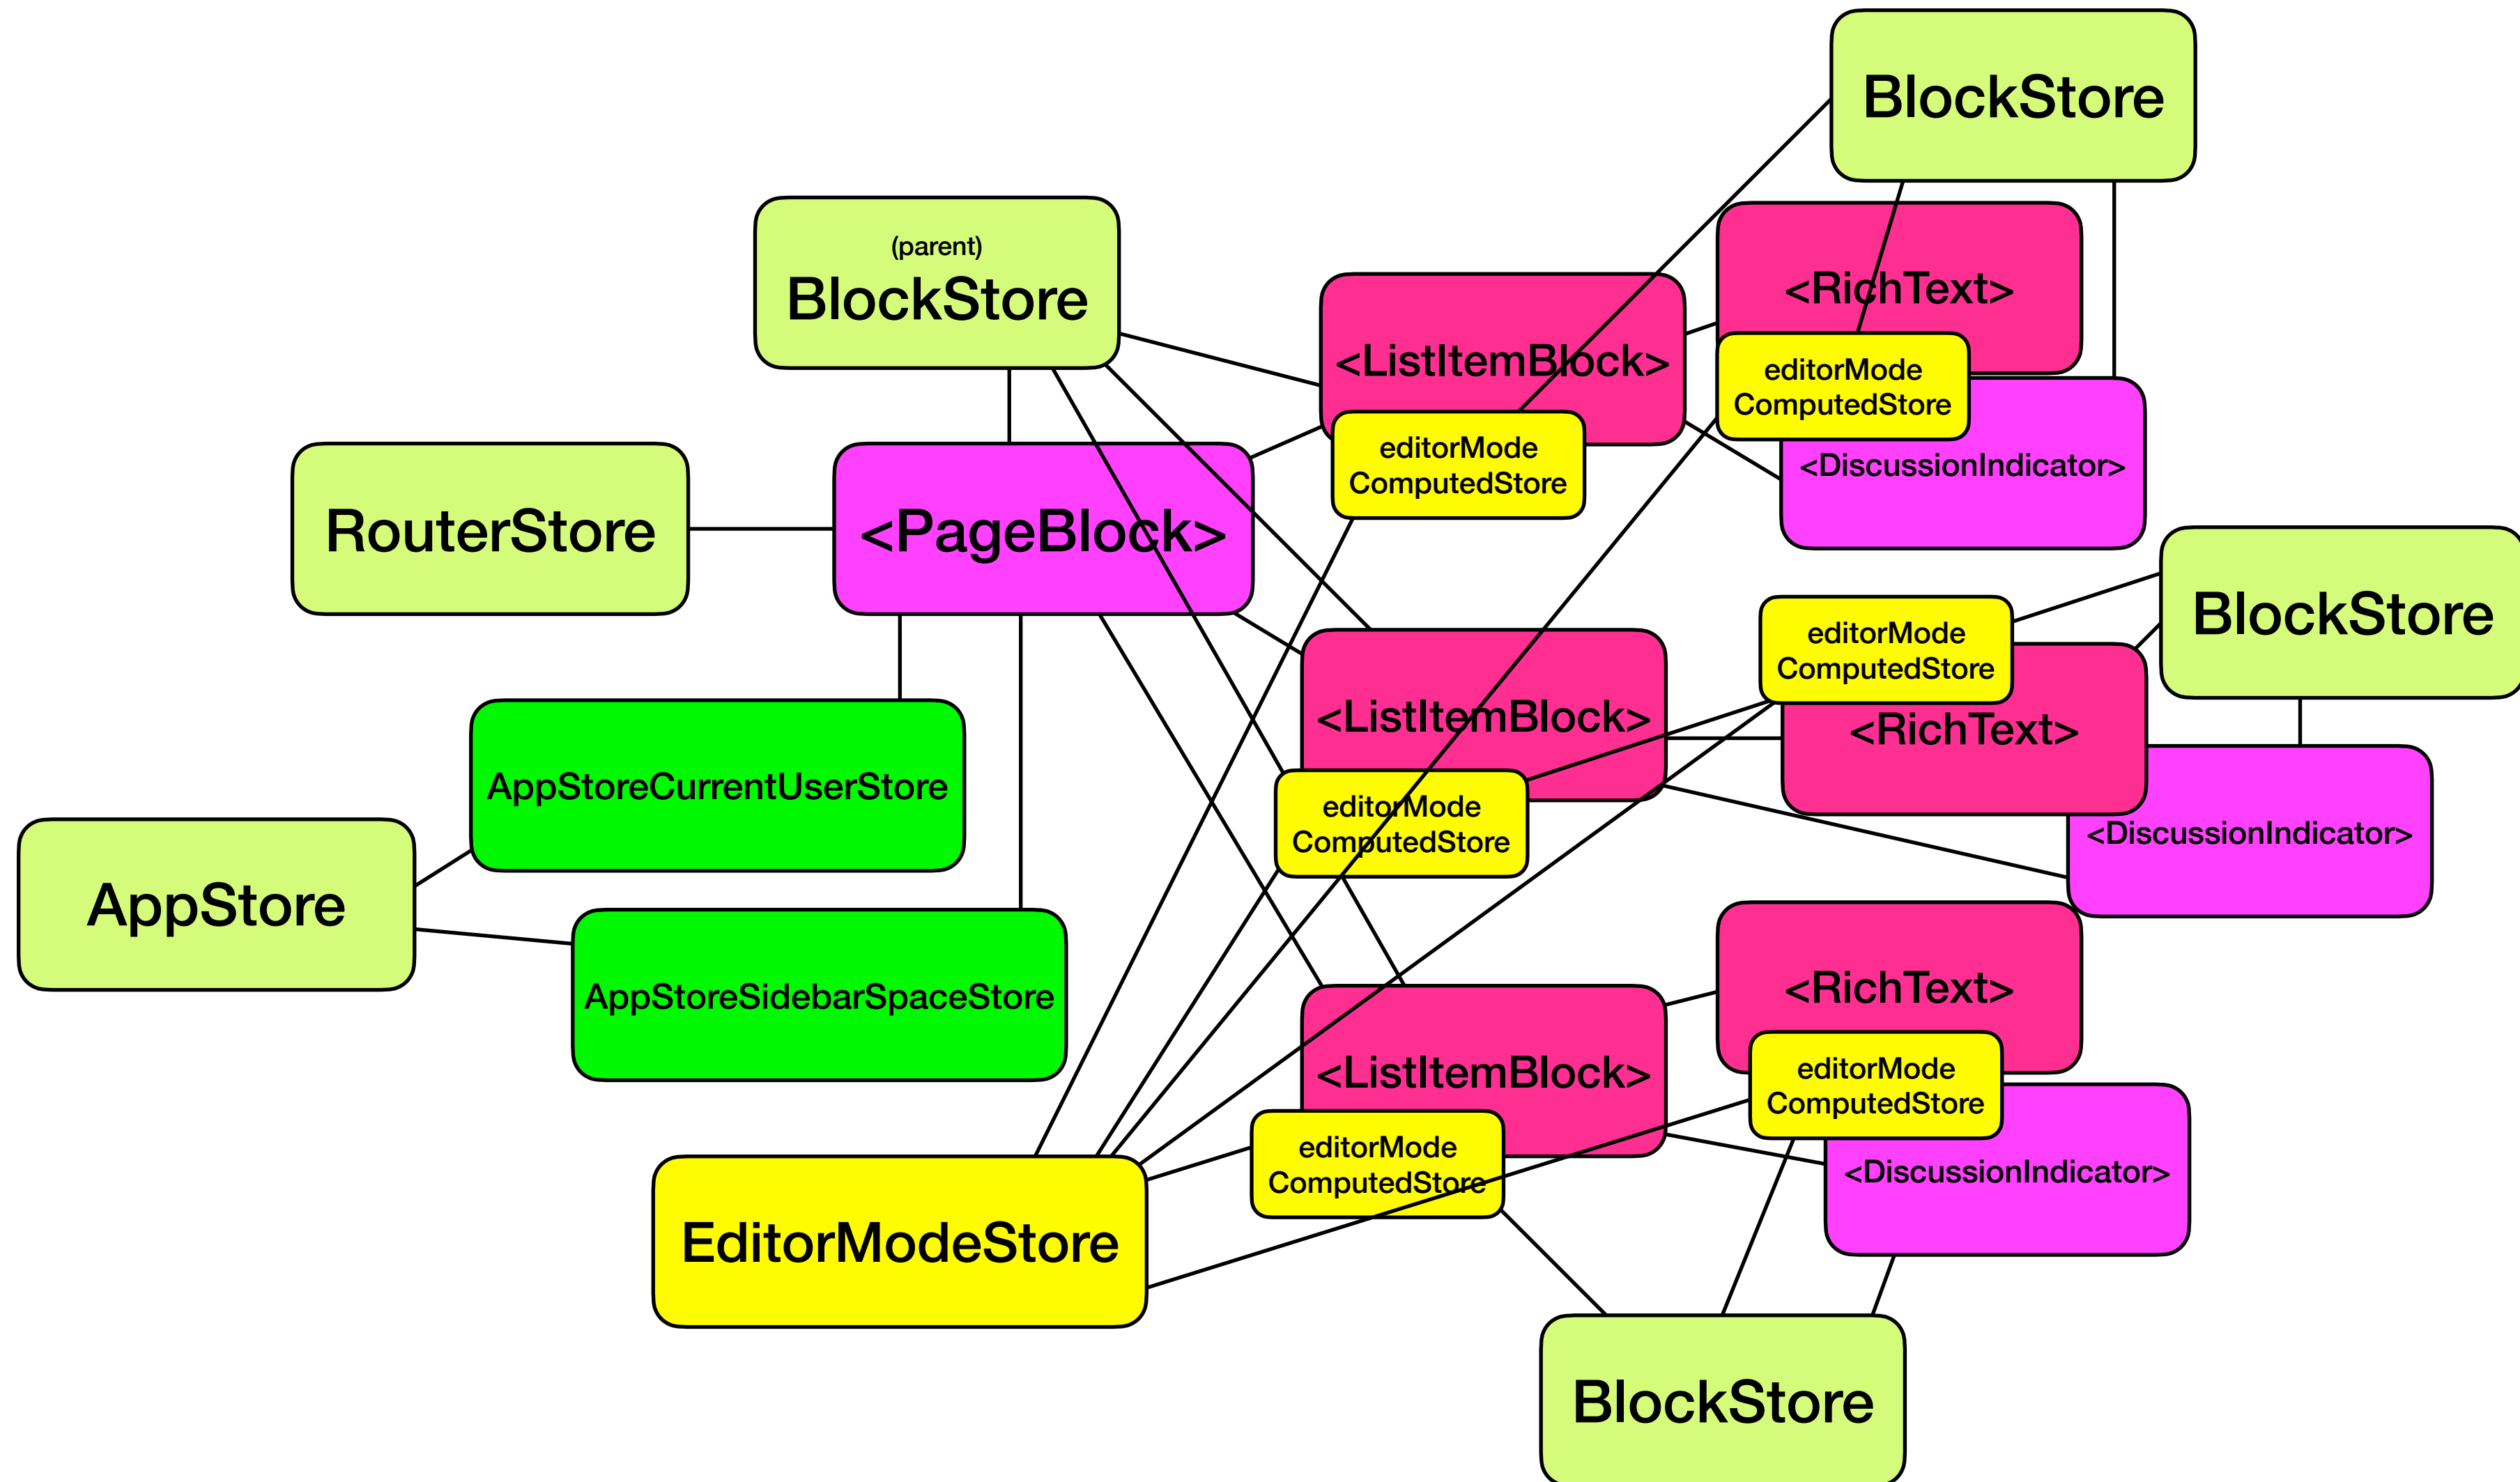
diagnosis: ~~murder~~ render



diagnosis: ~~murder~~ render



diagnosis: ~~murder~~ render



diagnosis: ~~murder~~ render

- what systemic patterns exist that are making the app slow?
 - unnecessary hook runs
 - unnecessary re-renders
 - unnecessary expensive hook wrappers

ai skills?

- translate the anti-patterns into a skill that can be run by engineers seeking to optimize their code
- speculative: can you run this skill in CI as a code review bot?

ai skills?

- want something generic? you might be interested in Vercel's react-best-practices or million.js's react.doctor

custom lint rules

- transform the anti-patterns into lint rules with the help of AI!
- the most surefire way to get an AI to follow instructions is to rely on engineering best practices

custom lint rules

- "hey claude, can you help me write a lint rule that ensures the return value of useComputedStore is a stable value? and can you make it autofix if possible? thanks!"

custom lint rules

- is this hook stable?
- does this need to be wrapped in this hook?
- is all of this prop used in the component?
- is this prop stable?
- is this hook duplicated within the component?
- should this computation be hoisted to a parent?
- are we only using a subsection of a returned value from a hook?

targeting problem

components

using chrome dev tools mcp

- "hey claude, can you load dev.annotation.com using Chrome Dev Tools MCP and click the AI button and profile how long it takes to open the chat popup?"
- honestly, this was disappointing!

using the profiler

- run a profile and ask the AI to find patterns in `*that*`
- "hey claude, I took this profile when clicking on the agent button. can you look at it and see what might be causing it to take the 600ms to open? thanks!"
- also works for the React Dev Tools profiler!

using the profiler

- "great! can you look at `Text.toHtml()` and tell me if you see any optimization opportunities?"
- AI is actually okay with vanilla JavaScript optimization!

speculative: heap snapshots

- this could also work for heap snapshots too, (especially if your company is tokenmaxxing!)

everything's so wrong
and I don't belong
living in your precious memory

you know i'd spend a hundred thousand tokens
if I could just see you!!

Constructor	Distance	Shallow Size	Retained Size
▶ (compiled code) ×614,084	3	49,784 kB 12 %	123,531 kB 30 %
▶ (string) ×307,917	3	32,877 kB 8 %	106,300 kB 26 %
▶ system / ExternalStringDa...	4	73,428 kB 18 %	73,428 kB 18 %
▶ (array) ×818,031	3	41,194 kB 10 %	70,939 kB 17 %
▶ I ×83,228	11	7,324 kB 2 %	56,440 kB 14 %
▶ Map ×158,918	4	2,543 kB 1 %	45,314 kB 11 %
▶ Object ×158,888	3	4,237 kB 1 %	44,819 kB 11 %
▶ IK ×38,666	6	3,867 kB 1 %	41,704 kB 10 %



diagnosing

regressions

OH NO



we're going up up up!!!!



using github mcp for regressions

- "hey claude, we saw a regression in clicking INP occur around June 18th. can you go through the PRs around that time and see if anything could have caused that regression? thanks!"

using AI tools for analyzing data

- "hey hexbot, can you make a graph of clicking INP by release version?"
- "hey hexbot, can you tell me if the INP regression was more prominent for certain cohorts, such as plan type, browser, or OS?"
- "hey hexbot, can you tell me if the INP regression was driven by any particular target?"

Lessons

Learned

Tried-and-true

engineering best

practices remain the

best practices

AI is still not good
at a lot of things!

AI tools should

supplement you,

not replace you!

thank you!

jenna.is/at-ny-perf-2026

@zeigenvector