

Twitter Thread by Corey Quinn



Corey Quinn

@QuinnyPig



Lest people think I've become too far lost in the Talking Head weeds this past #reinvent week, it's now time for me to livetweet my use of a new AWS service that was announced this week.

You probably get my re:Quinnvent newsletter. Here's the system I built that generates it.

Now that CodeGuru supports Python, let's turn it loose on the "build" subservice. This is a [@goserverless](#) stack. The subcomponent features 8 Lambdas function fronted by 11 APIs Gateway.

Yes, let's get started with it!

It starts with the correct option selected. Good for it, bad for CodeCommit's veneer of legitimacy.

Always always always use at least a 'project' or similar tag. Future you will thank you when you're attempting to figure out what the hell is costing money.

It's hanging out in "associating" state long enough that I'm starting to wonder if that's an ML term of art. GitHub's API is lightning-fast.

Bit of a UX fail here. To scan the entire repository I have to hit "code reviews" in the left nav bar. There's no direct shot from this window.

"Sure, I want security." Now after a bit of headscratching I see that it's overloading the term "repository" to mean an artifact bucket as well.

Much like a typical customer, I ignore that nonsense and disable the thing with the word "security" in it.

"master" is my feature branch because I was told to hold off on converting it to "main" until some GitHub tooling made this easier.

No, I'm not misusing the term "feature branch."

Much like I do, [@awscloud](#)'s CodeGuru names code reviews / commit messages by smacking a keyboard a few times. In this case, "nextweekinaws-master-6p394lx6vgx92drp"

We're now in state "pending." It claims to take 5-10 minutes.

Github tells me this app is 60.6% Python, 39.4% HTML.

98 commits, dating back to my initial commit in October of 2017.

And according to [@awscloud](#) CodeGuru Reviewer, my Python is perfect with no improvements possible.

This is of course a screaming red flag that this service is fake, or else my code is so bad that it's not even recognized as Python.

I have an unstaged commit; I'm ripping out legacy access tokens for [@Pinboard](#) and [@Pocket](#) because I don't use them anymore. The actual keys were stored in Parameter Store because at one point I liked to cosplay as a responsible engineer.

I create a pull request for that change, and find it's already analyzing it.

This is worrisome.

Because CodeGuru charges 75¢ per hundred lines of code in a pull request. Think of, say, massive version updates to packages, you add a new module, etc. These can get large--and you never had to think about it that way before. Now it's spendy.

And no recommendations to the pull request, which isn't a big surprise.

I add two more repos and start analyzing them.

(These are my custom anonymizing clicktracker for aggregate data only, and my "render Snarkdown into HTML" bits.)

Rude, [@awscloud](#).

I kid. It found no issues in any of my code. That's terrifying and now I know something's afoot.

Okay, *none* of my repositories are showing issues. I have a theory. Let's check the CodeGuru documentation.

Aha! Found it!

Remember, this is "Machine Learning," which is a polite euphemism for "bias laundering." Nobody knows how the algorithm works.

I suspect that it wound up looking at a bunch of code and making some assumptions, like "if the author is 'Stephen' the code

is shitty."

That is of course a joke. This is ML in tech in 2020. The actual author name would of course be something typically female-presenting.

So yeah. I don't honestly know how to evaluate this service. I assure you my code is far from perfect, but I can't seem to convince @awscloud of that.

I'll be expecting a job offer early next week.

One final note: if I weren't in the 90 day trial period, I would have paid @awscloud \$5 to tell me nothing at all via this experiment.

Oh, and even the pep8 linter has Some Thoughts on my shit tier code.