


# How much is that API call in the window?

Oct 16, 2023


Mark Robinson





# PLAID

- Plaid is focused on democratizing financial services through technology. We build beautiful consumer experiences, developer-friendly infrastructure, and intelligent tools that give everyone the ability to create amazing products that solve big problems.



Mark Robinson  
Infrastructure Engineer

---

**Do I need to do anything?**

---

The answer will shock and amaze

\$30/hr @ AWS

=

1 FTE

## What does a 10% reduction mean?

// 6

- If your monthly budget is above 6 figures, you should start investing in cost
- If you're above 7 figures, then you need to dedicate a team to this.
- Even if your spending isn't there yet, now is a great time to establish some good habits

---

## How does AWS make money?

---

Sing along if you know the words



## Super secret internal AWS memo

// 8

Me> How does AWS network transfer pricing work?

AWS> 

Charge 'em for the lice, extra for the mice

Two percent for looking in the mirror twice (Hand it over!)

Here a little slice, there a little cut

Three percent for sleeping with the window shut

When it comes to fixing prices

There are a lot of tricks I knows

How it all increases, all them bits and pieces

Jesus! It's amazing how it grows!



## What is AWS' value proposition?

// 9

- They sell reliability, performance, and scalability
- It's also very easy to spend
- You have real-time metrics for everything, except...
- Cost data is delayed by **8-24 hours** and has been broken for days at a time

---

## So how did we get here?

---

Don't blame me, I added blockchain support

## Debugging by adding resources

// 11

- One easy way to fix a broken service is to add more resources
  - Slow SQL -> bigger DB
  - Latency is high -> add CPU
  - System OOMs -> add memory
- It's not a bad idea at the start of your company's life
- As you grow you can't do this anymore

## Good code is cheap code

// 12

- Cost is a lagging indicator of quality
- If you go through your most expensive services you will find several serious problems
  - Unreliable
  - Hard to maintain
  - Slow to deploy
  - Unscalable

Paying more  
for worse  
software

---

## It's a sea change

---

Move under the sea, that's your solution to everything.

## It's an organization change

// 15

- Every engineer will need to care about costs a little bit
  - It can't just be one person doing this
- You will need forward looking tools
  - Performance debugging tools
- It's also a long process so pace yourself



## Do a Hackathon

// 16

- Find a few engineers who are interested and put them in a room for a few days
- Give it a fun name
- You'd be amazed at what you can find

## Create a cost experts group

// 17

- It's there to be consultants
- Help people understand the costs their architectures are imposing on you
- Don't be the fun police

## Build your cost analysis data

// 18

- You should start on this ASAP
  - Tagging is a great way to do this
- Decide what axes you want to assign costs
  - You should have a cardinality of 10s-100s
- Make the right thing, the easy thing
  - Make this settings required in all your configs

---

## Scaling up

---

Boy, that escalated quickly.

# Insights

// 20

- Develop your insights enough to make your next decision and stop (for now)
- Focus on building durable insights rather than one offs. The durable insights will pay off later on
- Establish policies that enforce ownership

## Decisions

// 21

- Once you have enough information, you will have a list of possible improvements
- You don't need perfect information
- Don't let people argue out of improvements because "we're too small"

## Action

// 22

- Pick fixes that are easy to implement first
  - Cross team dependencies are hard
  - Customer affecting fixes are death
- Don't be seduced by going after single large fixes
- Repeated small fixes everywhere can bring large cost savings at low risk
- Don't ignore small easy wins because all dollars are equal



---

## No one can intuit what things cost

---

It's one server, Michael. How much could it cost? \$10/hr?

## Fix your RFC process

// 24

- Make people write up an estimate for how much their new feature will cost
- The AWS cost calculator is pretty bad
- Make your own in excel
- Focus on easily understood costs and then add a “tax” for indirect costs

# Write a cost calculator

// 25

6							
7	Service	Count	Unit rate	Line Total			
8	Pod count		--	--	Sub Total	\$0.00	
9	CPU Request (millis) (from the scaffold file requests sections)		\$0.04	\$0.00	Overhead (This is networking costs, logging volume, metrics, and other AWS fees)	\$0.00	
10							
11	S3 Storage (GiB)		\$0.02	\$0.00	Monthly Estimate	\$0.00	
12	S3 Glacier (GiB)		\$0.004	\$0.00			
13							
14	EC2 Instances (excl k8s) (m5.large equivalents)		\$0.096	\$0.00			
15	EC2 Instances (excl k8s) (c5.large equivalents)		\$0.085	\$0.00			
16	EC2 Instances (excl k8s) (r5.large equivalents)		\$0.126	\$0.00			
17	EBS Disk (GiB)		\$0.15	\$0.00			
18							
19	MySQL Storage (GiB)		\$0.15	\$0.00			
20	MySQL instances (r5.large equivalents)		\$0.29	\$0.00			
21							
22	Redis Instances (m6g.large equivalents)		\$0.15	\$0.00			
23							
24	Mongo Storage (GiB)		\$0.50	\$0.00			
25							

## Write up a list of possible fixes

// 26

- Your teams won't know where to start
- Write up a doc on things they can do to fix things
- Customize it per-team with cost estimates
  - People LOVE this
- It's super easy because all cost problems are pretty similar at a high level

---

## Inevitable problems

---

GASP!

## What could go wrong

// 28

- Cost blow outs
  - Security scanners
  - Sudden scale-up in requests
  - Log volume spikes for debugging

# Take Aways

// 29

- All dollars are the same
  - Start with the easy wins
  - Make friends and find allies
- Start your cost insights monitoring
  - Today
    - Now
- Good software is cheap software
  - How much is a good night's sleep worth?
- No one knows what they're doing
  - Give engineers the tools to understand the problem



# Thank you

Come find me later to talk how about how  
you're overpaying for S3