

Tips and Tricks for Building a Scalable Cloud Service

HackArizona 2015

Alex Landau
Amazon

Amazon

- Office in downtown Tempe near ASU (boooooooooo)
- Most teams work for **Marketplace** (third party sellers)
- ~80 developers and growing fast
- **We're hiring.**



The cunning plan...

- What? Why?
- Learn by doing
- Implications

What does it mean to be Scalable?

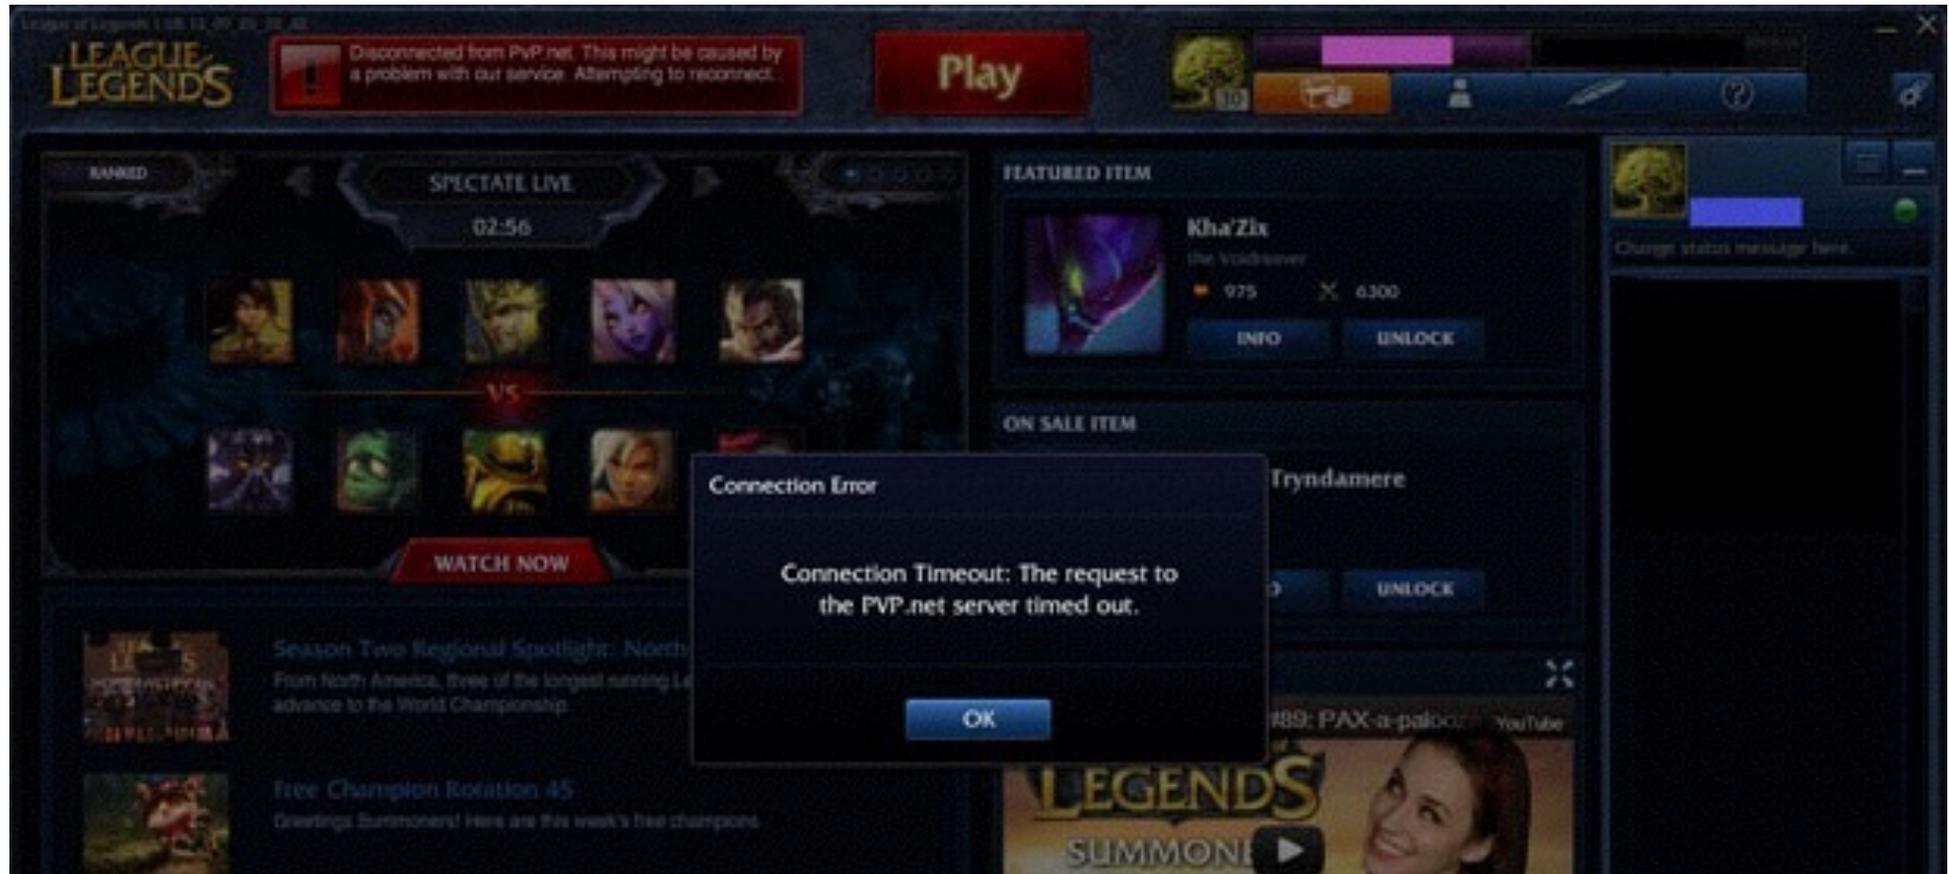
- Available



reddit is under heavy load right now, sorry. Try again in a few minutes.

What does it mean to be Scalable?

- Available
- Really, really fast



What does it mean to be Scalable?

- Available
- Really, really fast
- **Reliable**



What does it mean to be Scalable?

- Available
- Really, really fast
- Reliable

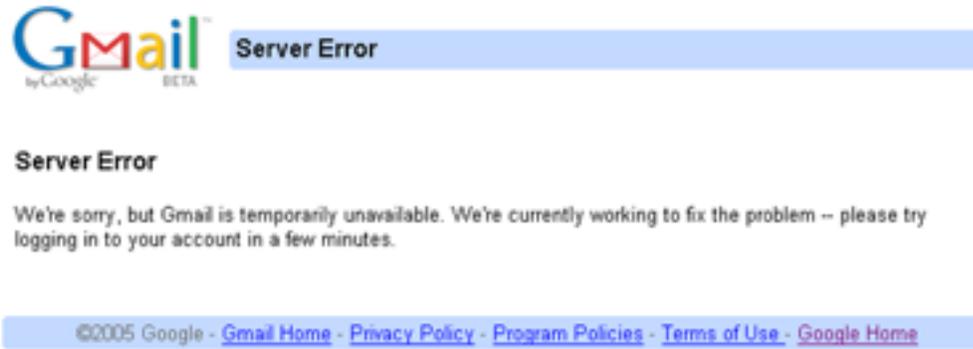
- **With lots and lots of traffic.**

What is a lot of traffic?



- LA has ~374,000 cars per day on I-405 (<http://tinyurl.com/adlsuol>)
- Reddit has ~25 million upvotes and downvotes per day (<http://www.reddit.com/about/>)
- Dropbox hit 1 billion file uploads per day in 2013 (<http://tinyurl.com/kskhn28>)
- Amazon's S3 service peaks at 1.5 million requests **per second**. (<http://tinyurl.com/q6hnqvl>)

Why do we care?



Gmail by Google BETA

Server Error

Server Error

We're sorry, but Gmail is temporarily unavailable. We're currently working to fix the problem -- please try logging in to your account in a few minutes.

©2005 Google - [Gmail Home](#) - [Privacy Policy](#) - [Program Policies](#) - [Terms of Use](#) - [Google Home](#)



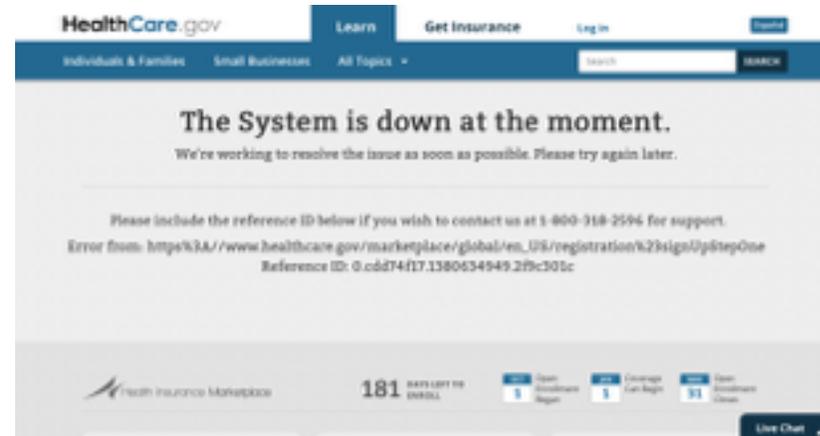
looks like we shouldn't have stopped using itp...

- **User experience should be positive. These things suck.**



The top screenshot shows a message: "Hey, your account is temporarily unavailable due to site maintenance. It should be available again within a few hours. We apologize for the inconvenience." Below it are login fields and a "Login" button.

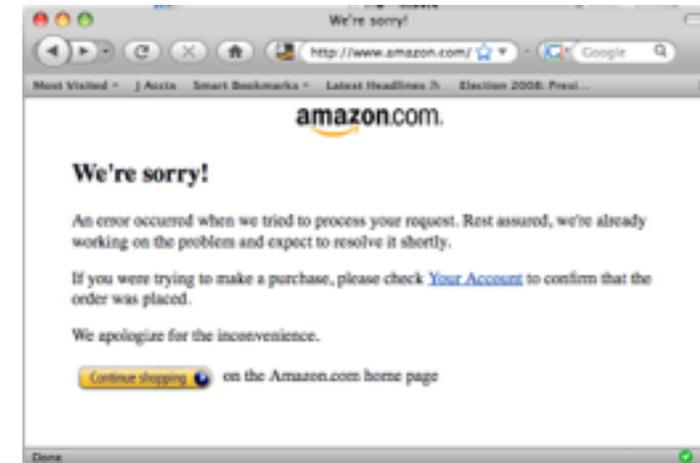
The bottom screenshot shows a message: "Sorry, an error has occurred. We're working on getting the feed as soon as we can."



The System is down at the moment.

We're working to resolve the issue as soon as possible. Please try again later.

Please include the reference ID below if you wish to contact us at 1-800-318-2596 for support.
Error from: https://www.healthcare.gov/marketplace/global/en_US/registration/k2/signUpstepOne
Reference ID: 0.cdd74f17.1380634949.2b0301c



We're sorry!

An error occurred when we tried to process your request. Rest assured, we're already working on the problem and expect to resolve it shortly.

If you were trying to make a purchase, please check [Your Account](#) to confirm that the order was placed.

We apologize for the inconvenience.

[Continue shopping](#) on the Amazon.com home page



Digg has broken an axle.

We might have to sell some oxen but we'll be back on the trail soon.



WIKIMEDIA FOUNDATION

Fout Fel Follo 错误 错误 Erreur Error Fehler エラー Blad Errore Erro

The Wikimedia Foundation servers are currently experiencing technical difficulties. The problem is most likely temporary and will hopefully be fixed soon. Please [check back](#) in a few minutes. For further information, you can visit the [#wikipedia](#) channel on the [Freenode](#) IRC network. In the meantime, you may be able to view [Google's cached version of this page](#). Wikipedia is now one of the [most visited](#) sites on the Internet by traffic and continues to grow, and as a result the Wikime has a constant need to purchase new hardware. If you would like to help, please [donate](#).



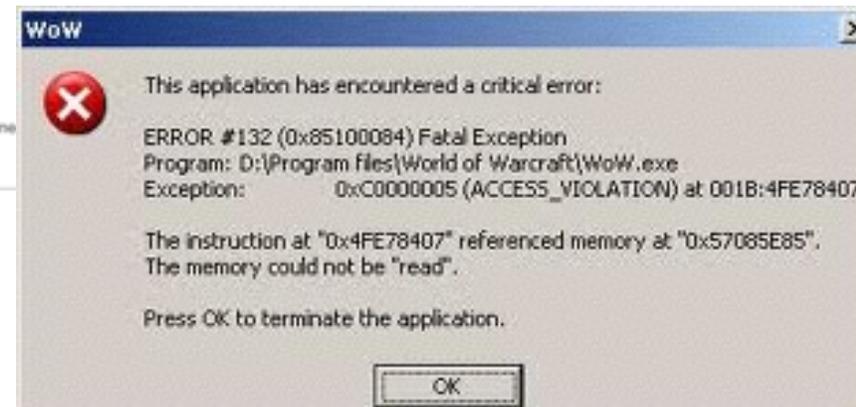
Help empower the world with free knowledge!

One time gift of \$ (USD)

Public comment (200 characters max)

Public donor list List my name List anonymously

[Donate Now!](#)



Wow

This application has encountered a critical error:

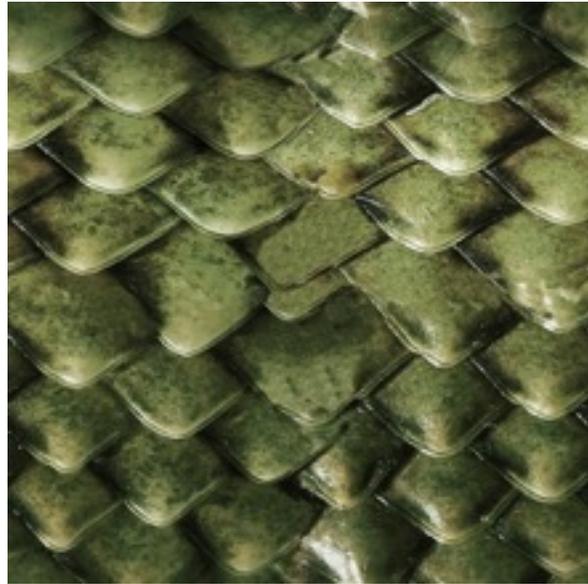
ERROR #132 (0x85100084) Fatal Exception
Program: D:\Program files\World of Warcraft\Wow.exe
Exception: 0xC0000005 (ACCESS_VIOLATION) at 001B:4FE78407

The instruction at "0x4FE78407" referenced memory at "0x57085E85". The memory could not be "read".

Press OK to terminate the application.

OK

So how do we build a service that scales?



note names: C D E F G A B C D^b E^b F G^b A^b B^b C D^b

D E F[#] G A B C[#] D E^b F G A^b B^b C D E^b

E F[#] G[#] A B C[#] D[#] E F G A B^b C D E F

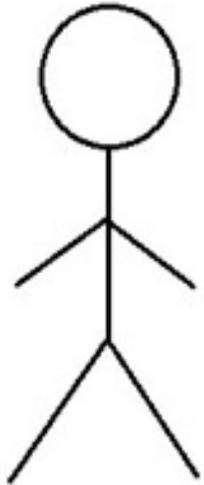
F[#] G[#] A[#] B C[#] D[#] E[#] F[#] G A B C D E F[#] G

A^b B^b C D^b E^b F G A^b A B C[#] D E F[#] G[#] A

B^b C D E^b F G A B^b B C[#] D[#] E F[#] G[#] A[#] B

The basics – our little service

User



ID, name,

Location
n



ID, name, GPS coords,
...

Checkin
n



User was at location
at timestamp.

The basics – our little service

User		
ID	Name	...

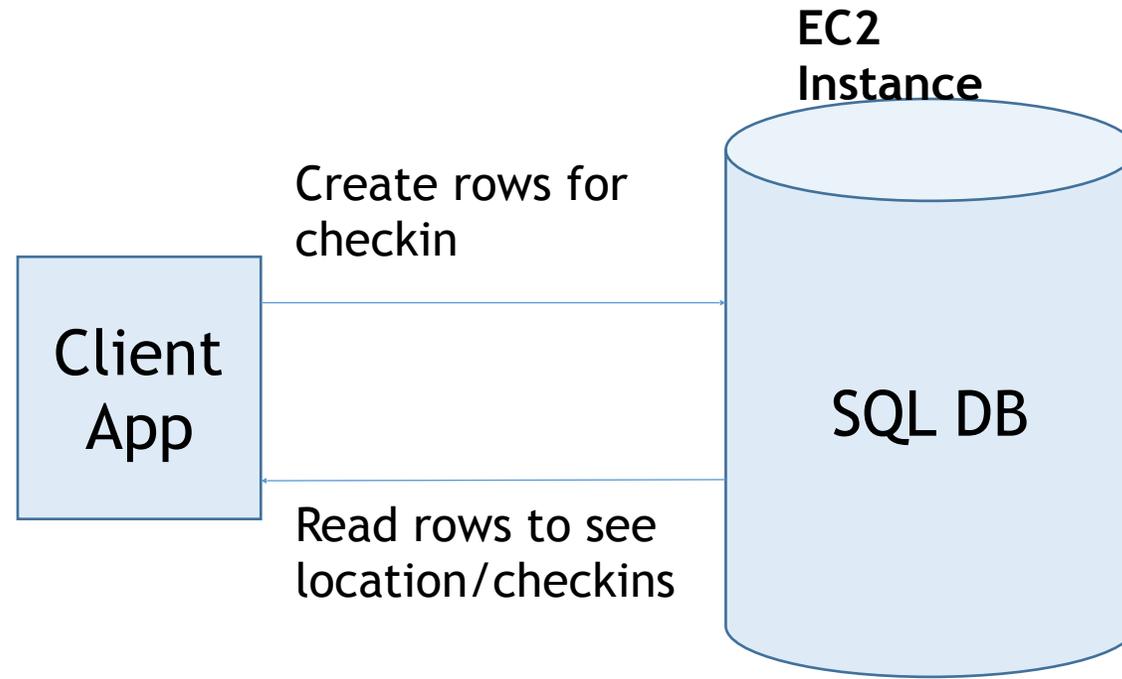
Location			
ID	Name	GPS Coords	...

Checkin			
ID	User ID	Location ID	Timestamp

The basics – our little service

- Record a checkin
- See a user's current location
- See the checkins for a user across a time range
- See the checkins at a location for a user across a time range

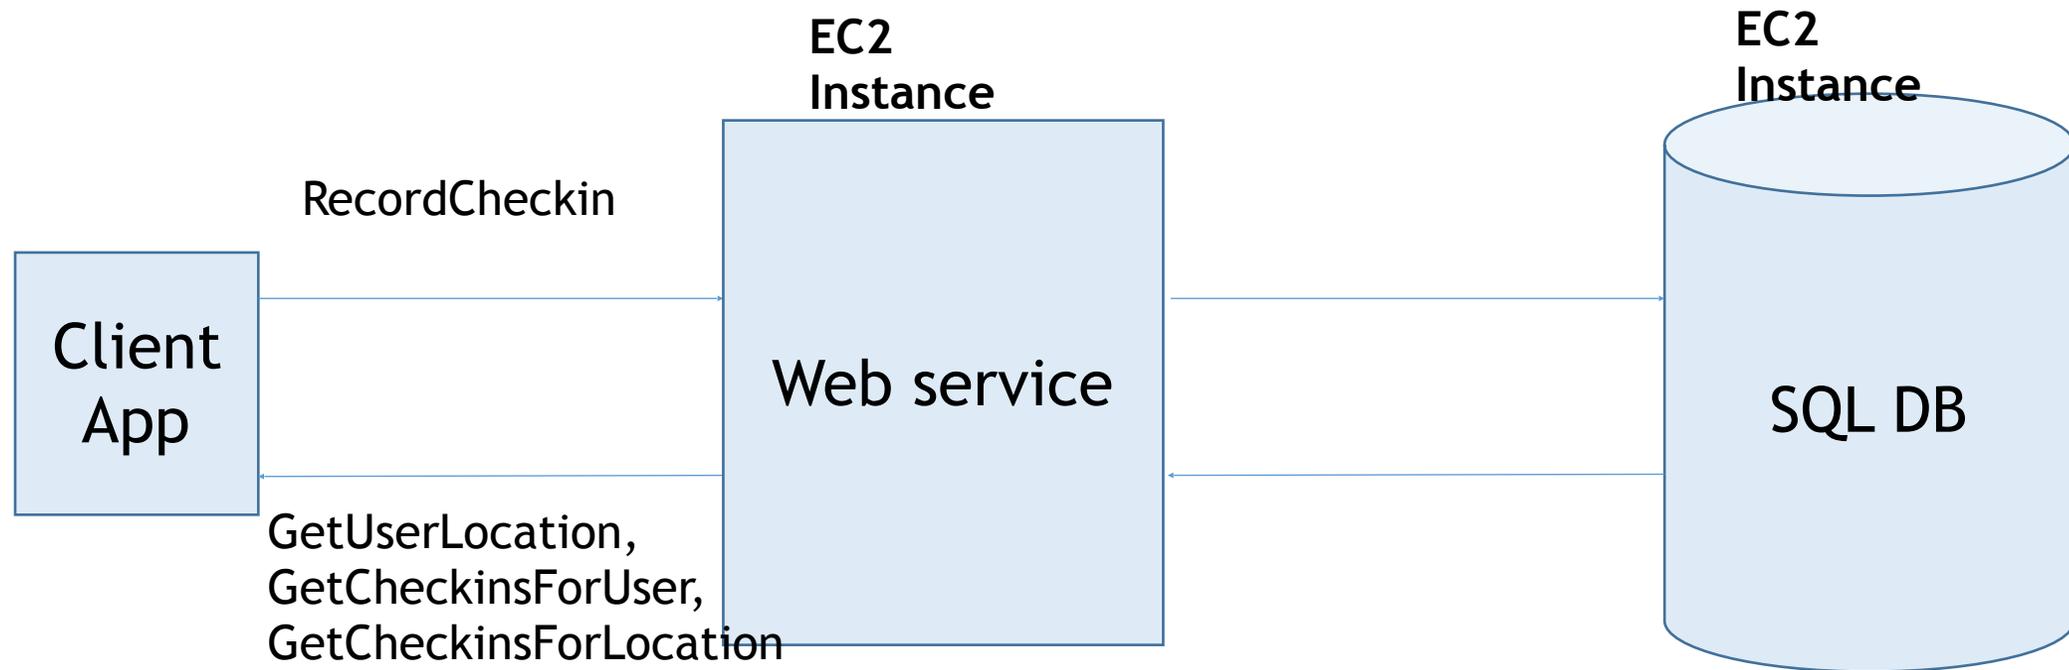
Attempt 1 – Direct database access



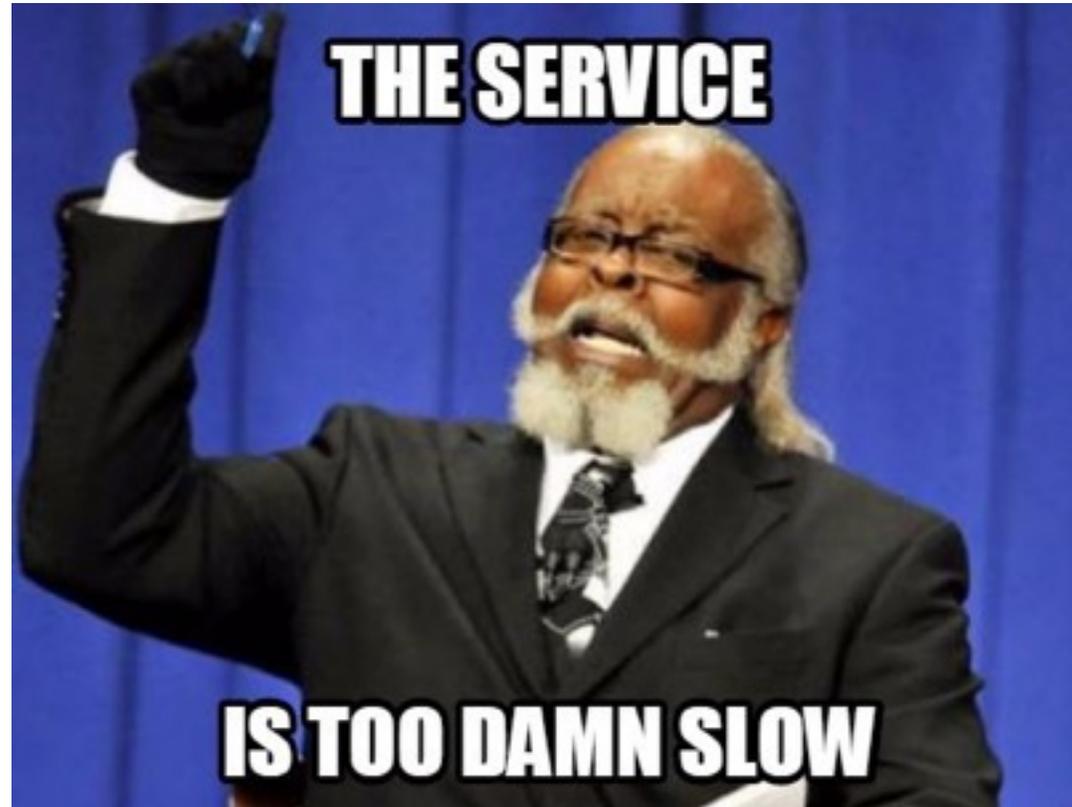
Problems



Attempt 2 – A service middleware



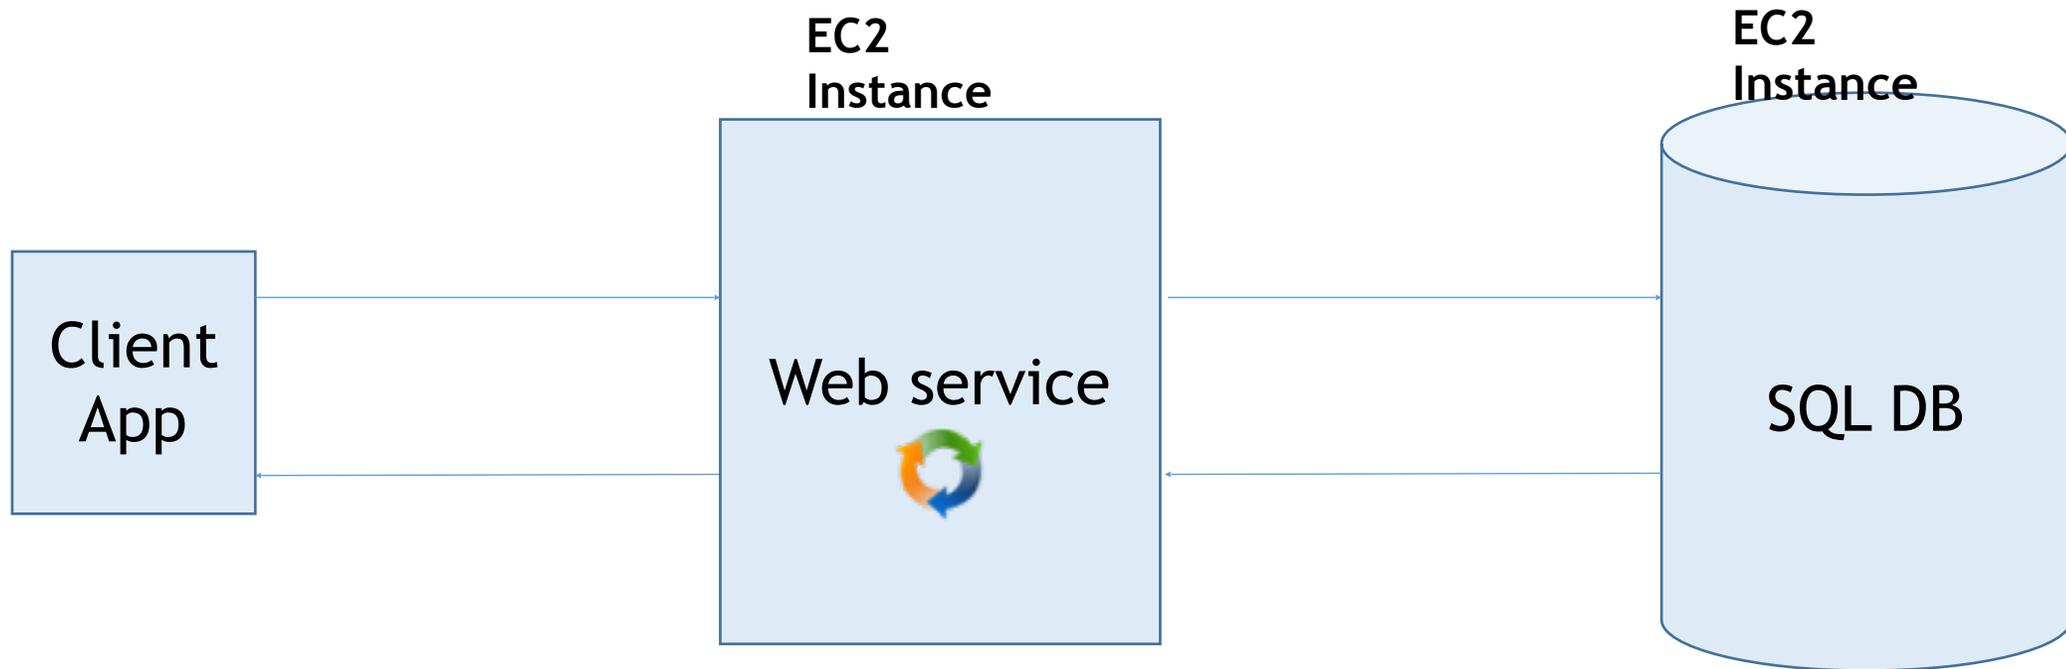
Problems



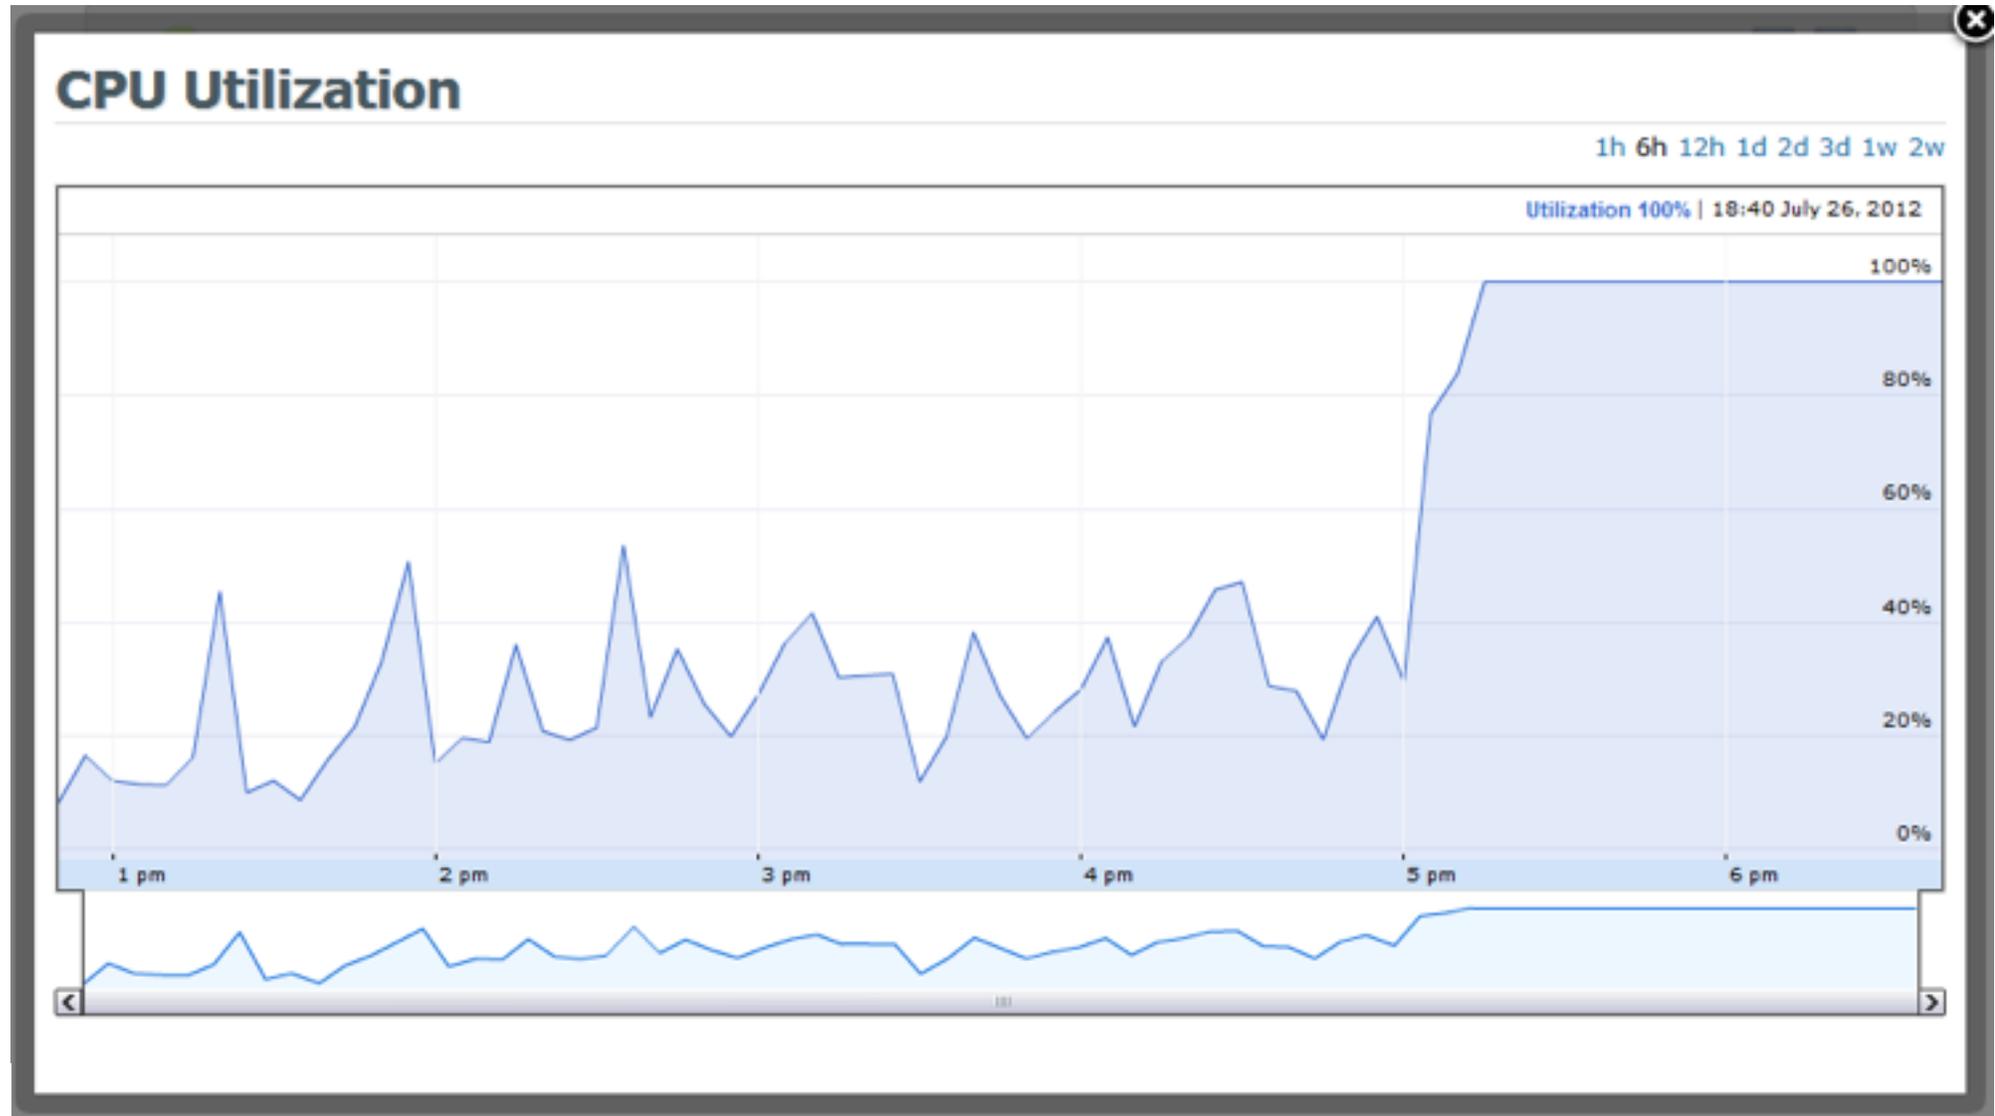
Asynchronous IO



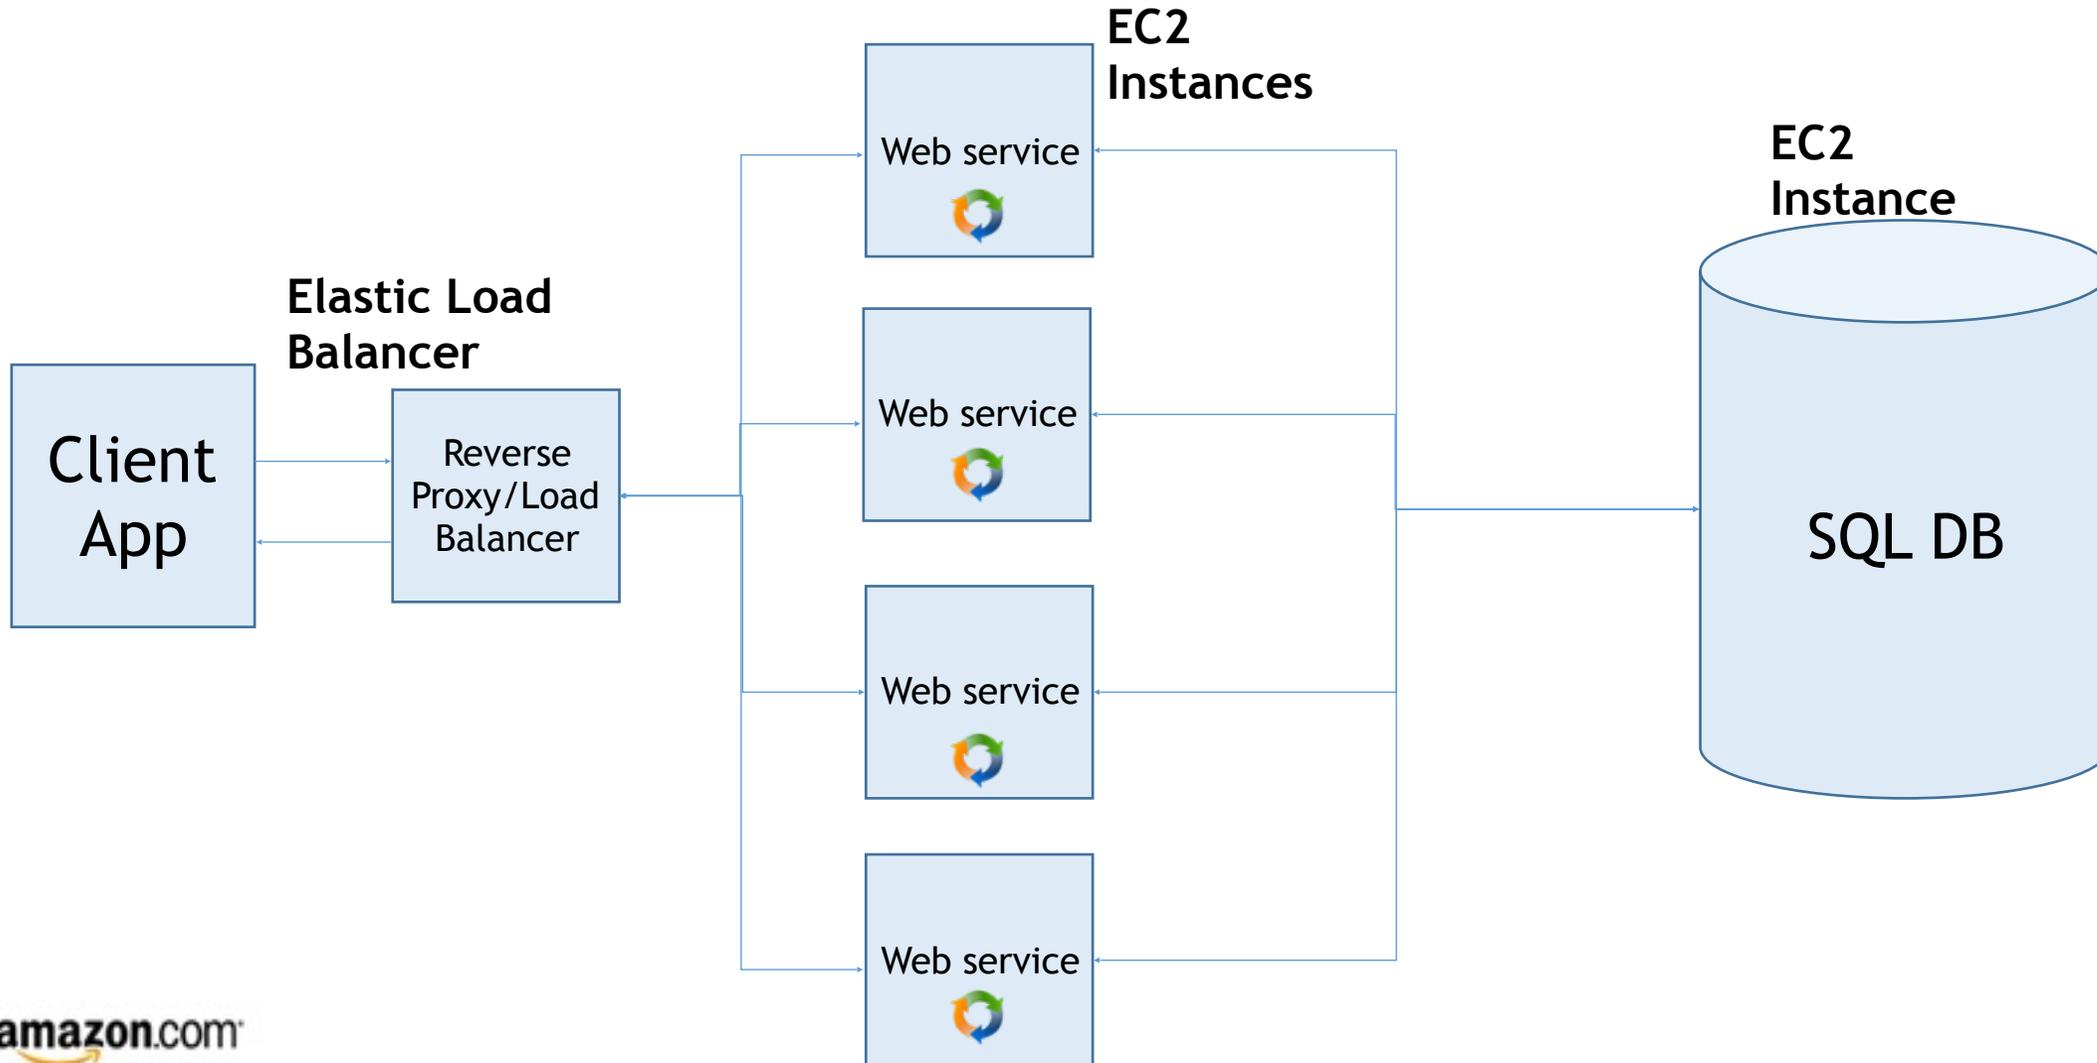
Attempt 3 – Asynchronous IO



Problems

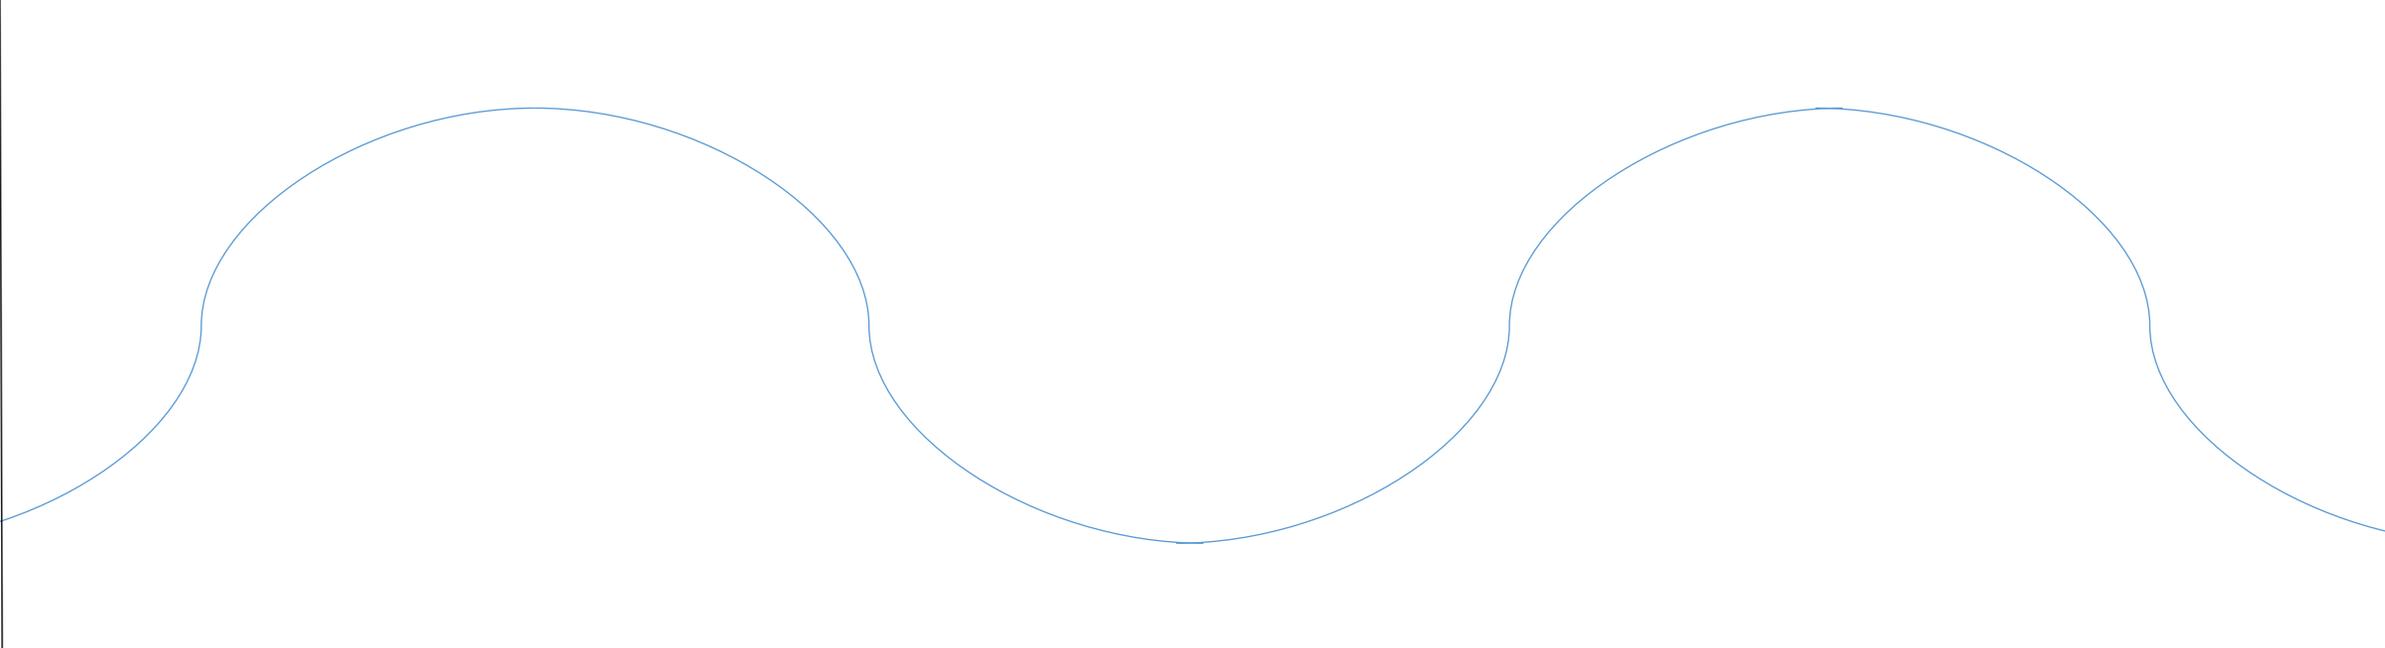


Attempt 4 – Load balancing



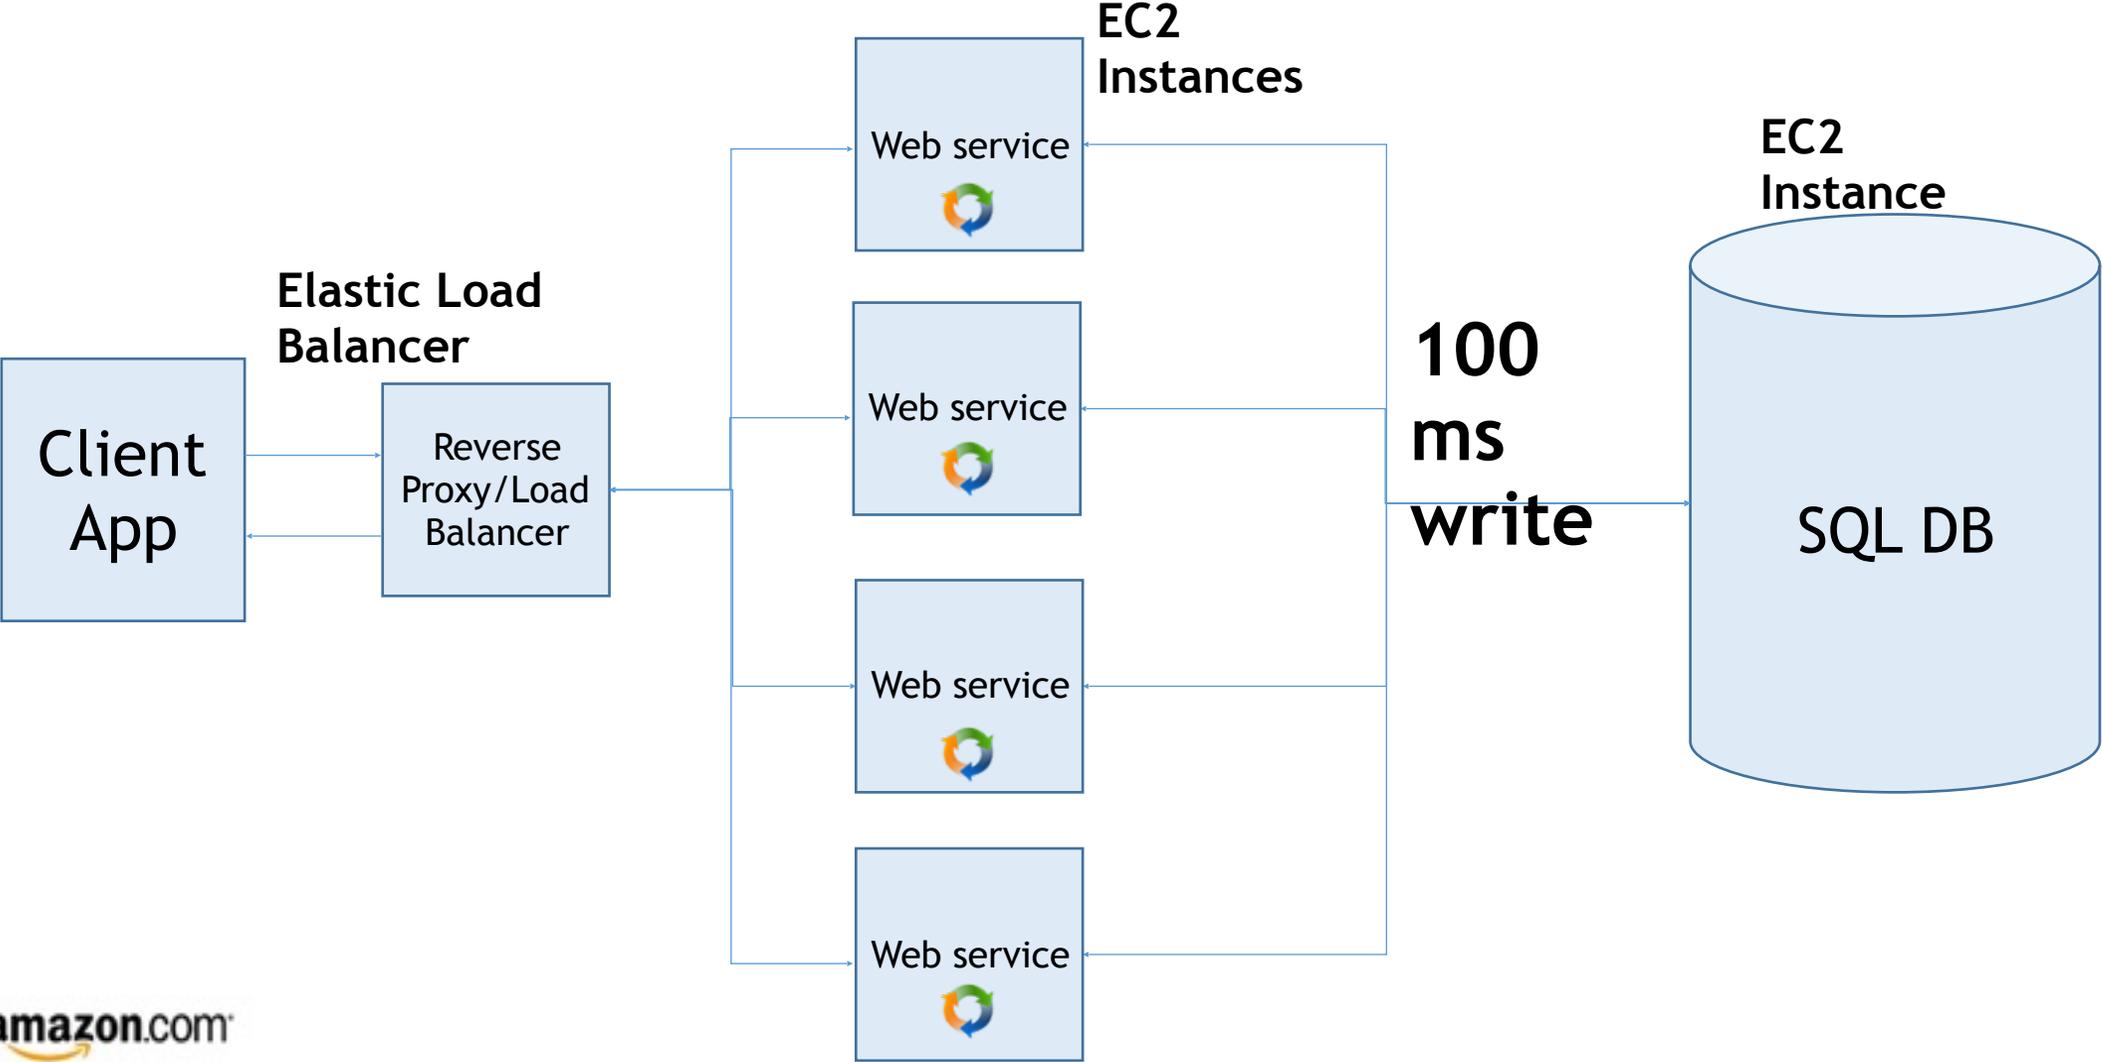
Problems

Database Writes

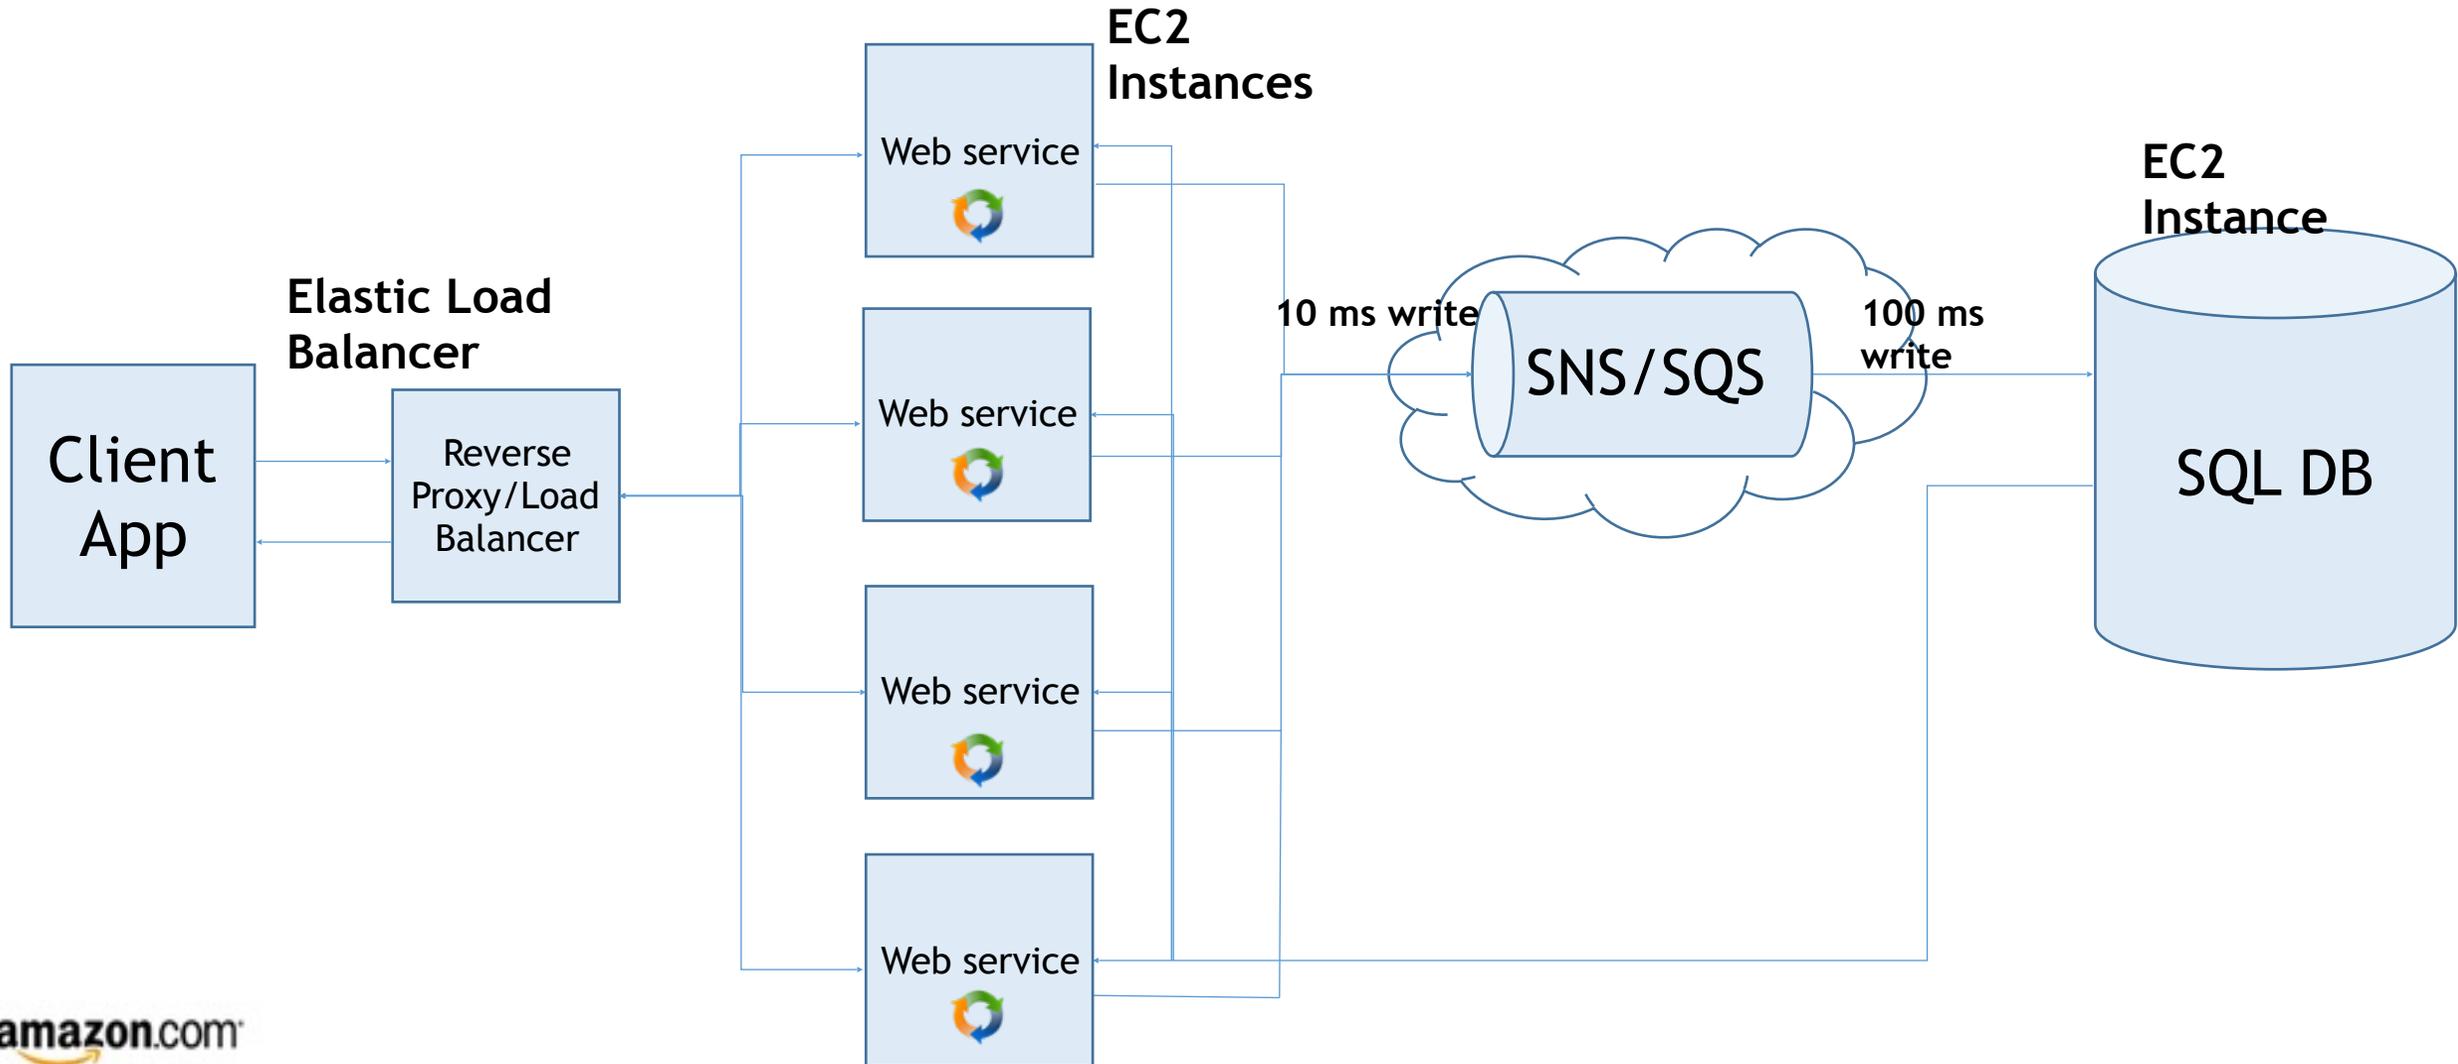


Time

Problems



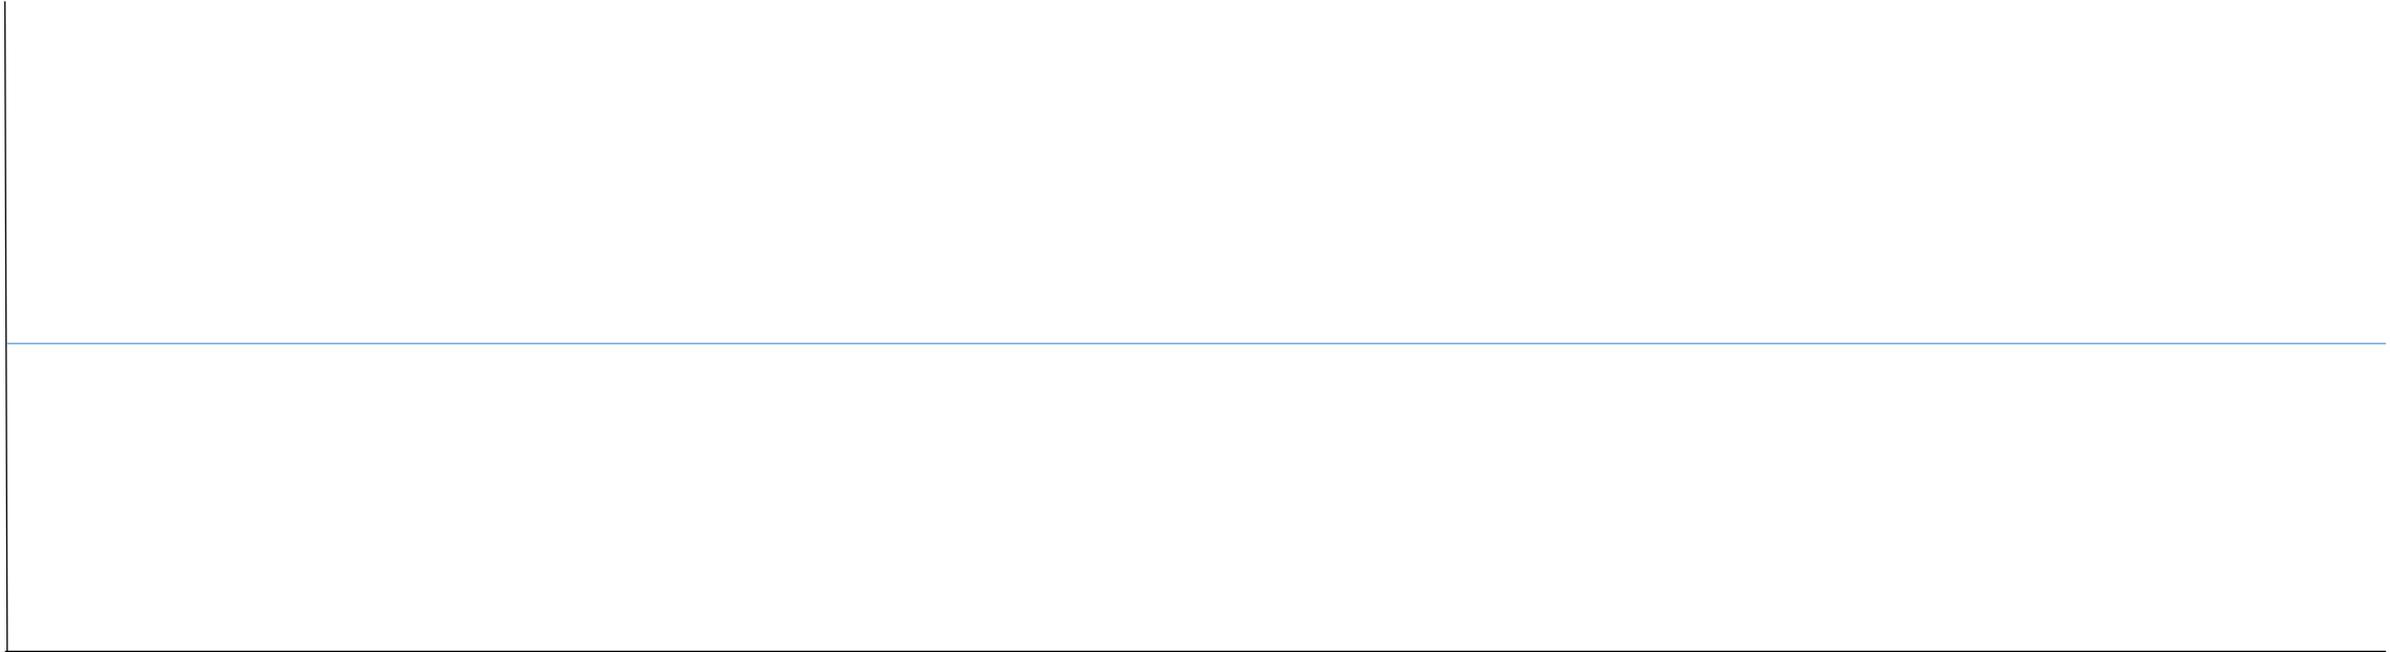
Attempt 5 – A queue in the cloud



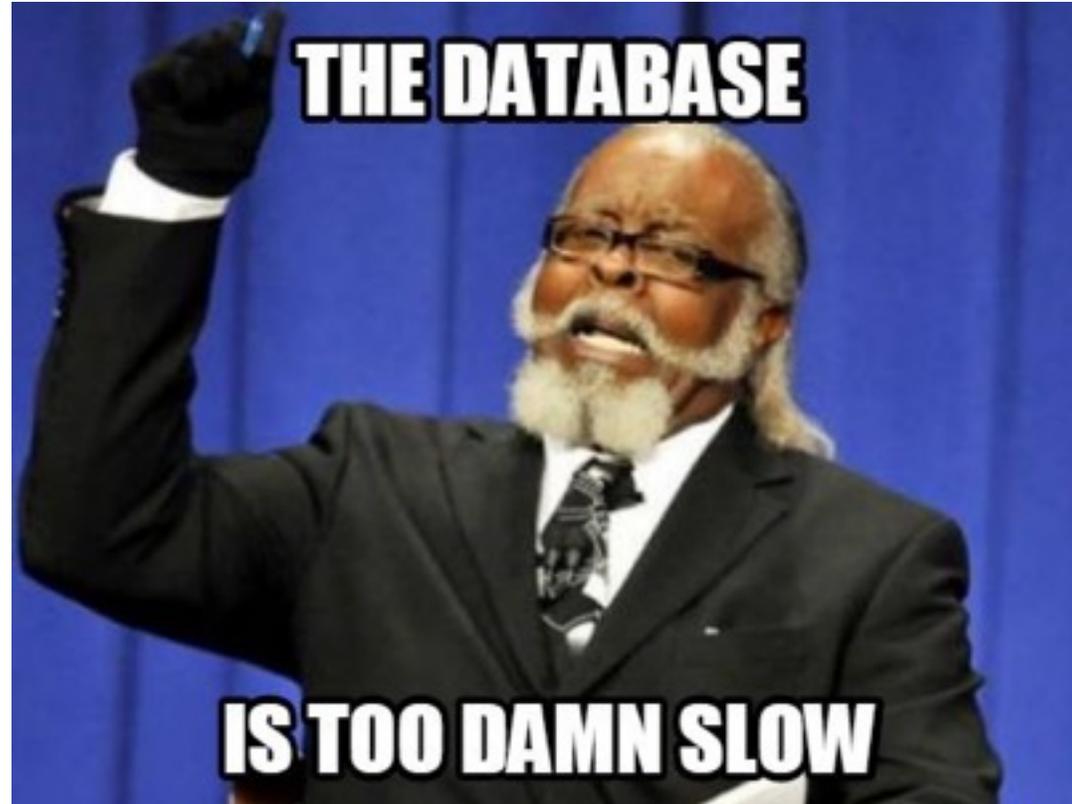
A queue in the cloud

Database Writes

Time



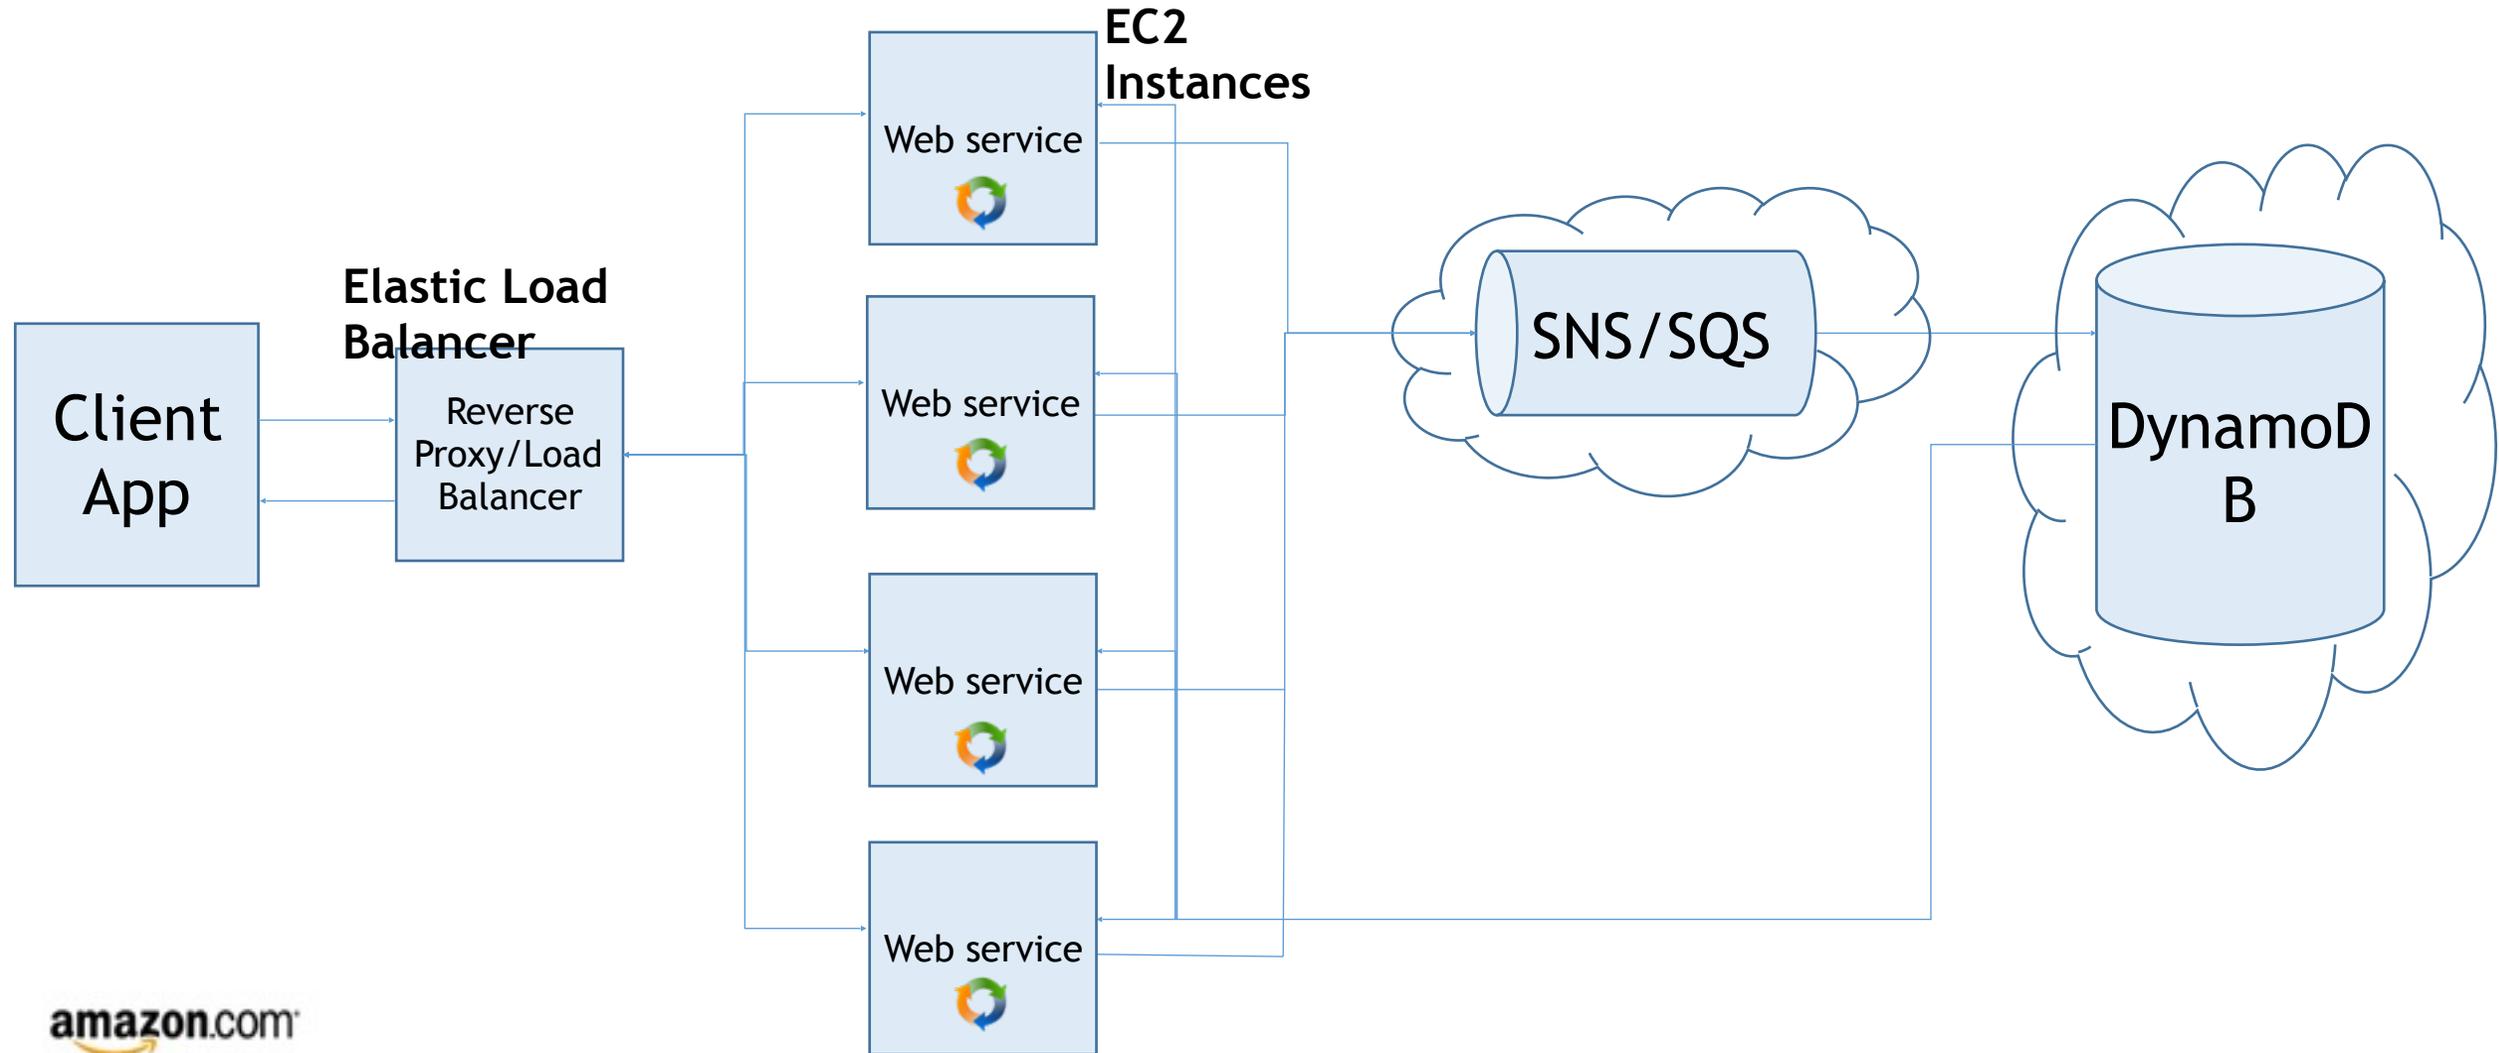
Problems



Amazon DynamoDB

- NoSQL
- **Fast, available, reliable**
- Automatically scalable
- ↑ Throughput == ↑ \$\$\$

Attempt 6 – A database in the cloud



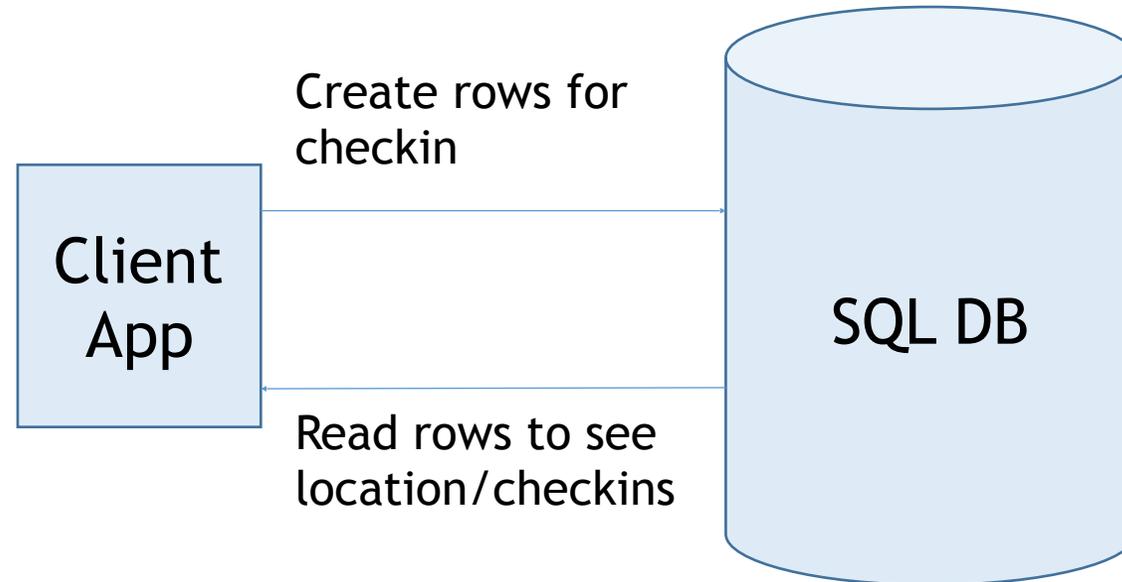
Even more improvements

- **CloudSearch** to speed up checkin reads
- Time-series database (**Druid, InfluxDB, Prometheus**) for checkins
- Friends lists, push notifications, metrics and data
- **Caching (ElastiCache)**



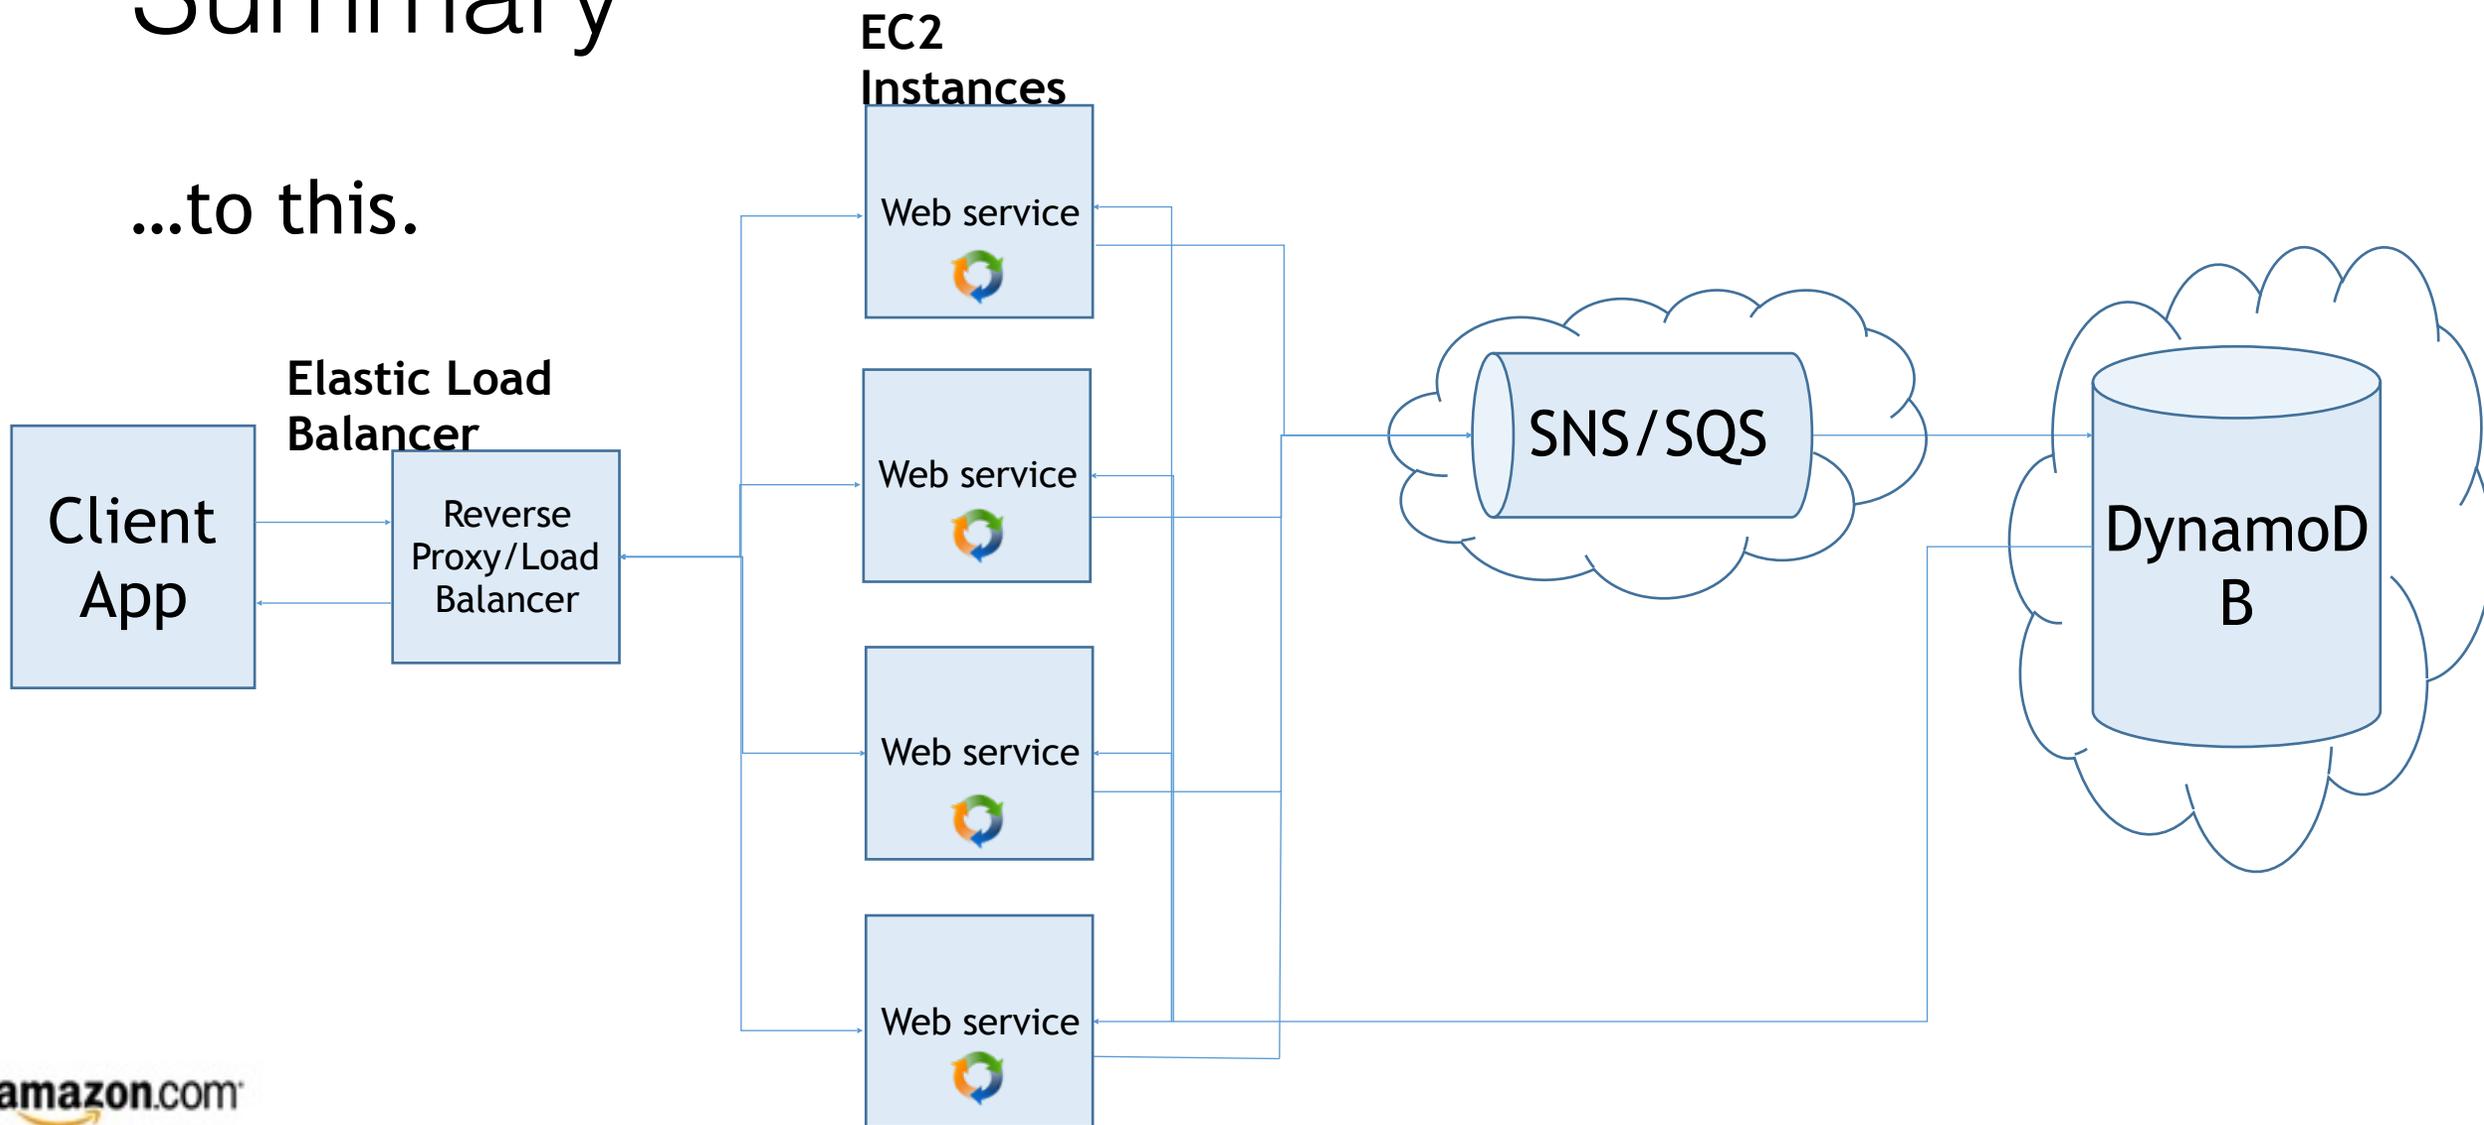
Summary

We went from this...



Summary

...to this.



Summary

- Big wins.
 - Linearly scalable.
 - More money → handle more traffic.
 - Availability, performance, reliability.
 - Managed services ≈≈ **significant** savings. (<http://tinyurl.com/q29n6h3>)

Implications

- Real world services are significantly more complex, **but they use the same techniques described here.**
 - Asynchronous IO
 - Messages queues
 - Managed cloud databases (both SQL and NoSQL) with service guarantees
 - Load balancing
 - Caching
- You can make your entire infrastructure completely **horizontally scalable.**

Questions?

