

The logo for Cloud Day 2022 features a stylized cloud shape on the left, with three vertical lines of varying heights extending downwards from its base, resembling a circuit board or server rack. To the right of this icon, the text "CLOUD DAY 2022" is written in a bold, black, sans-serif font.

CLOUD DAY 2022

**Evoluzione di una piattaforma cloud managed
in un'applicazione serverless moderna**

Saverio Tosi
Spreaker from iHeart
 **@SaverioTosi**



SPONSOR

DIAMOND



GUCCI



PLATINUM



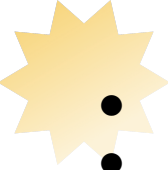
GOLD



PARTNER




Evoluzione di una piattaforma cloud managed in un'applicazione serverless moderna

- 
- Spreaker is a **podcast platform**
 - Within **Spreaker**, podcasts can be hosted, distributed, monetized, discovered and listened to
 - Part of **iHeartMedia** family since 2020

Spreaker 
From  **iHeart**





- Spreaker is a podcast platform
- Within Spreaker, podcasts can be hosted, **distributed, monetized**, discovered and listened to
- Part of iHeartMedia family since 2020



Data from April 2022

**400 Millions
IAB Certified
Downloads**

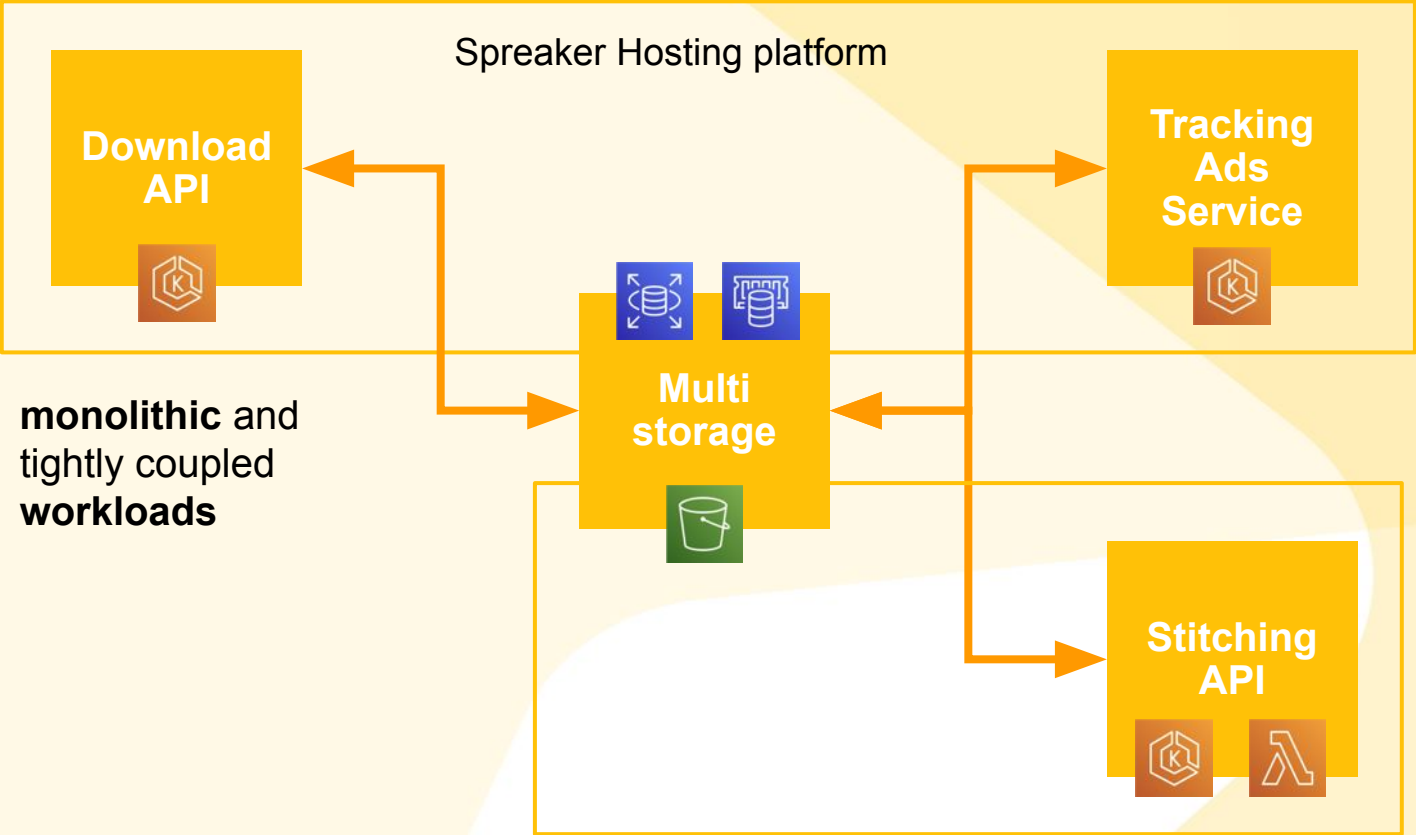
**25 PB of audio
delivered through
Cloudfront**

Podcast Delivery Network

The Big Rewrite

Why

Podcast Delivery Network in 2021



Multi Storage



What do we need to store?



Media

With media we mean the file audio of a podcast and its metadata.



Variant

A variant is a combination of Ads and one episode.

4.5M variants per day



Session

It is a download from a specific user (IP+UA) of a specific episode (variant)

avg 22M sessions

Multi Storage in 2021

- RDS Aurora Postgres “media-ads”
- S3 buckets
- Redis “Cache”
- Redis “Session-state”



What's wrong?



What were the options?

We were looking into consolidating as much as possible into a single storage. Without increasing cost and ruining user experience.




Postgres

- Expertise in Spreaker
- SQL Database;



Redis

- Fast;
- Key/value database
- Automatically delete old items;



DynamoDB

- Fast;
- Flexible NoSQL database;
- Automatically delete old items;
- Serverless;

Which one should we use?



What were the options?

We were looking into consolidating as much as possible into a single storage. Without increasing cost and ruining user experience.




~~Postgres~~

- Expertise in Spreaker
- SQL Database;



~~Redis~~

- Fast;
- Key/value database
- Automatically delete old items;



~~DynamoDB~~

- Fast;
- Flexible NoSQL database;
- Automatically delete old items;
- Serverless;

Can we use S3?

PRO:

- Scale automatically
- Lifecycle events
- “Cheap”

CONS:

- “Performance”

We ran some test to understand how S3 would perform and it turned out that...

Yes, we can!

avg 50ms

UPLOAD

avg 25ms

GET

Note: Head 15ms



Simple Storage Service



S3 - Optimizing Performance

- Amazon S3 automatically scales to high request rates
 - Your application can achieve at least 3,500 PUT/COPY/POST/DELETE and 5,500 GET/HEAD requests per second per **prefix** in a bucket.
 - There are **no limits** to the number of prefixes in a bucket.
 - Example:
 - *variant/1234-567-89.json* WRONG
 - *variant/1234/567/89.json* CORRECT

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/optimizing-performance.html>



S3 - Optimizing Performance

- Make sure your AWS SDK reuses existing connection.
 - AWS SDK for Javascript v2 has this option **disabled** by default
 - AWS SDK for Javascript v3 has this option **enabled** by default
 - You can enable it using the env var `AWS_NODEJS_CONNECTION_REUSE_ENABLED`

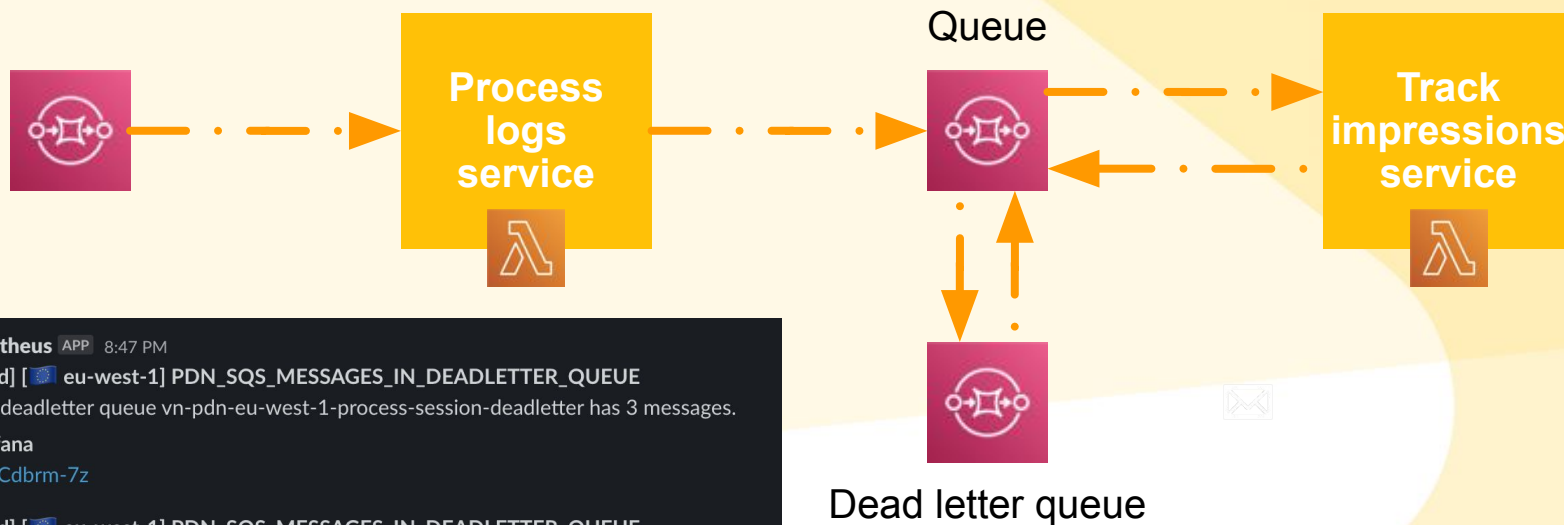
<https://docs.aws.amazon.com/AmazonS3/latest/userguide/optimizing-performance.html>

~~— Multi Storage~~
- Decoupling

Event driven application: Tracking ads service



Decoupling, handle failures and data recovery



Prometheus APP 8:47 PM


[prod] [eu-west-1] PDN_SQS_MESSAGES_IN_DEADLETTER_QUEUE
The deadletter queue vn-pdn-eu-west-1-process-session-deadletter has 3 messages.

Grafana
[d/NCdbm-7z](#)

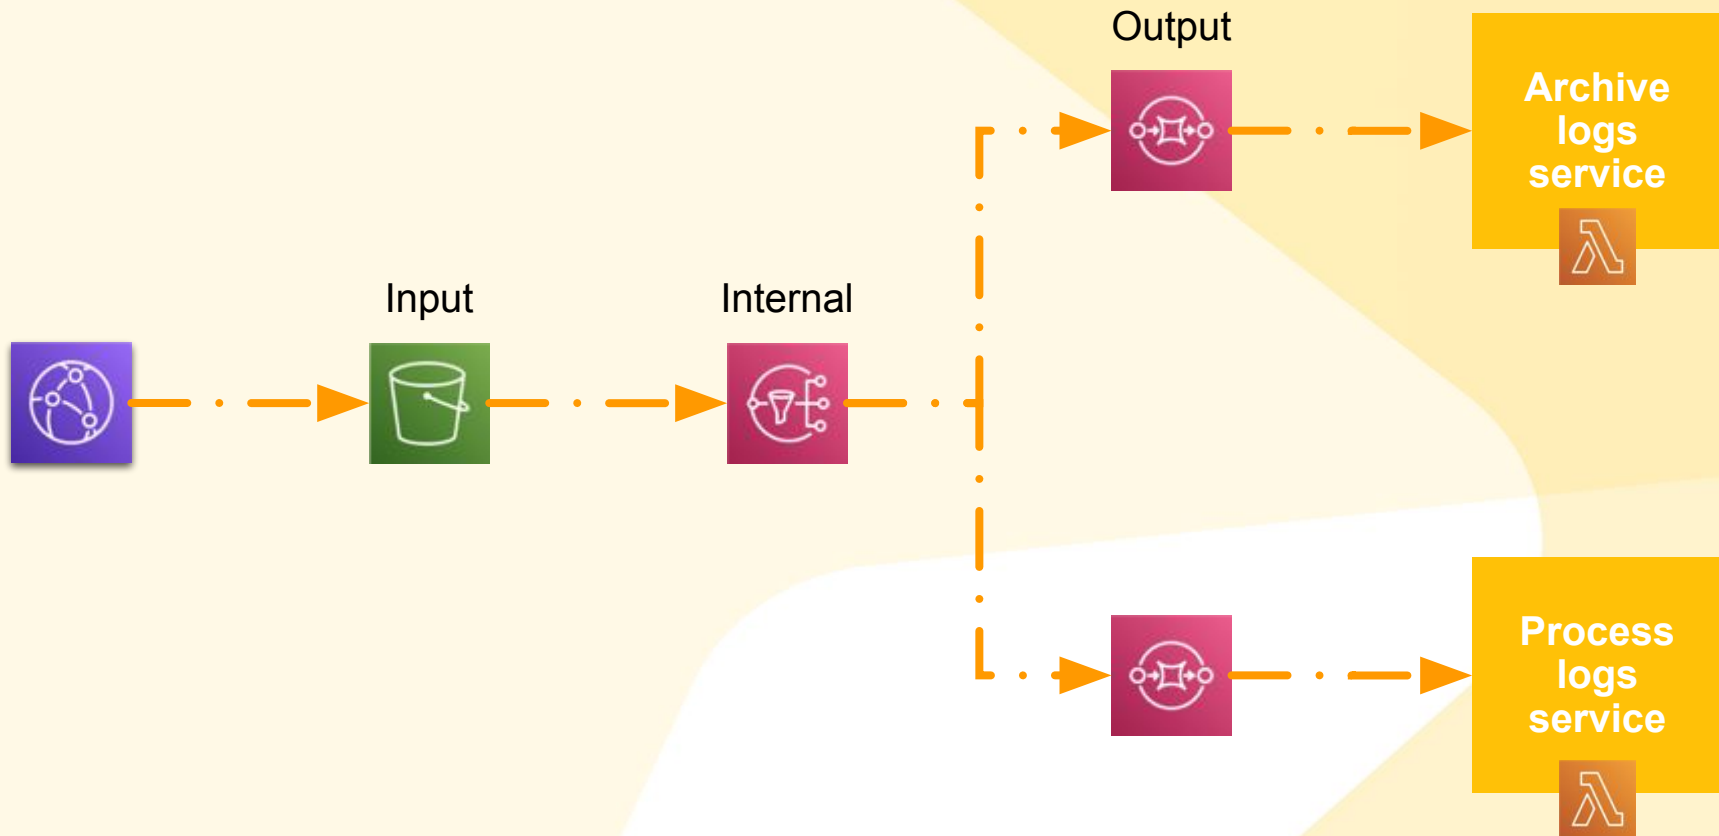
[prod] [eu-west-1] PDN_SQS_MESSAGES_IN_DEADLETTER_QUEUE
The deadletter queue vn-pdn-eu-west-1-track-impression-deadletter has 950 messages.

Grafana
[d/NCdbm-7z](#)

Scalability and reduce costs with batches

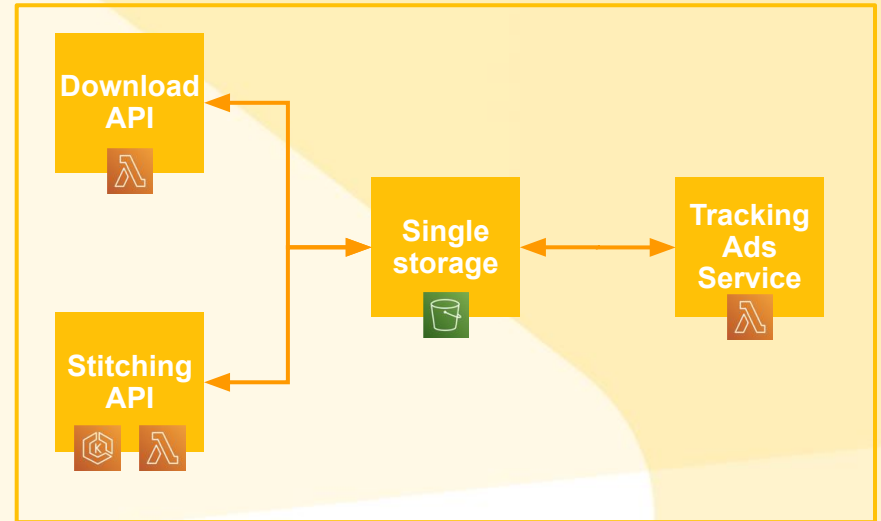
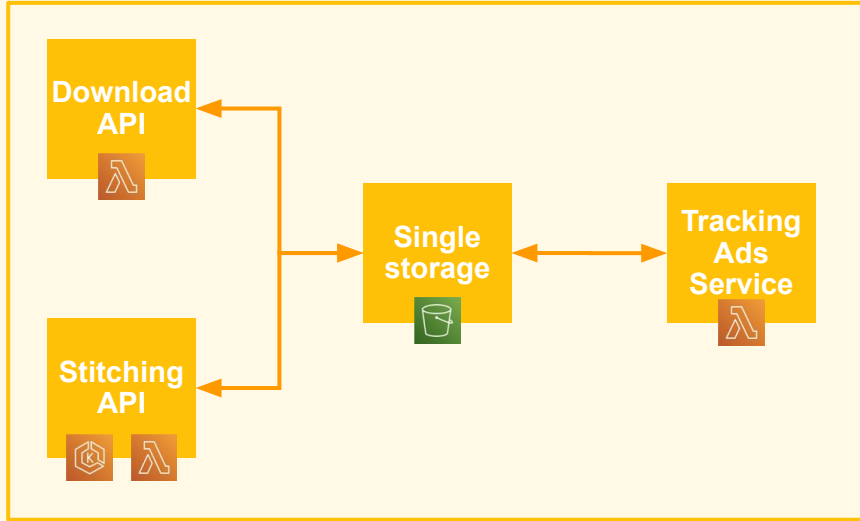


One single event, multiple triggers



~~— Multi Storage~~
~~- Decoupling~~

Podcast Delivery Network in 2022





Podcast Delivery Network in 2022

Reduced TCO

We no longer have to:

- upgrade/apply security patch to our db
- scale up or out db and ec2 instances

Scale 10x the traffic

We removed the main bottleneck of the previous implementation.

Single storage

Thanks to S3 we got

- high availability
- scalability
- multi-region
- retention
- event driven



Conclusions

- Can we use serverless?
 - yes, we can. but...
- What is the best ...?
 - It depends

We are hiring

<https://careers.spreaker.com/>

Grazie!