



PROCESSING DATA LAKE WITH NODE.JS IN SERVERLESS ARCHITECTURE





Nikolay Matvienko


Software Engineer at Grid Dynamics

You can find me at twitter.com/matvi3nko

github.com/matvi3nko



DATA LAKE




RAW DATA MODEL

STRUCTURED DATA

ANALYZED DATA MODEL



THE GOAL



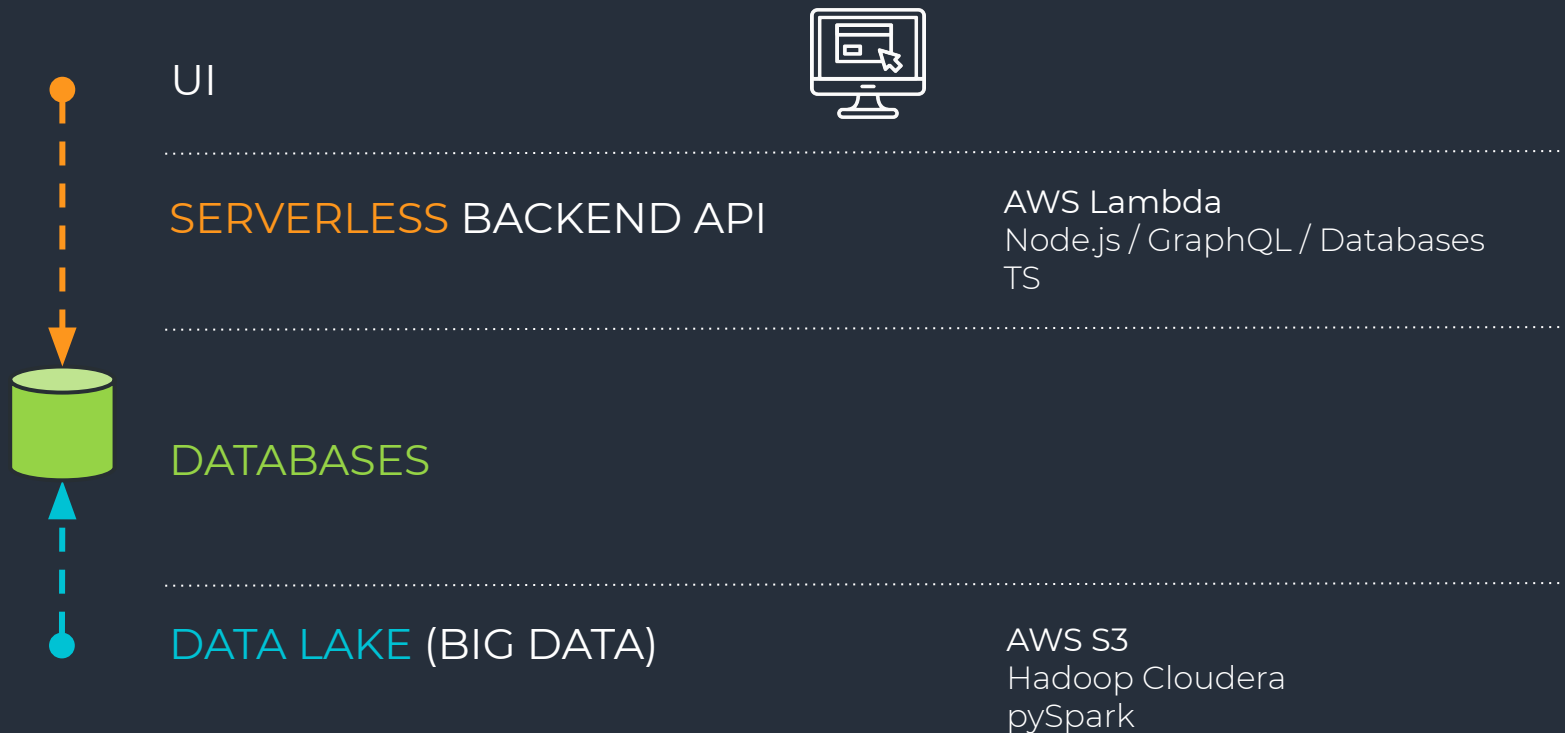
CASE STUDY

INTELLIGENCE SUPPLY CHAIN

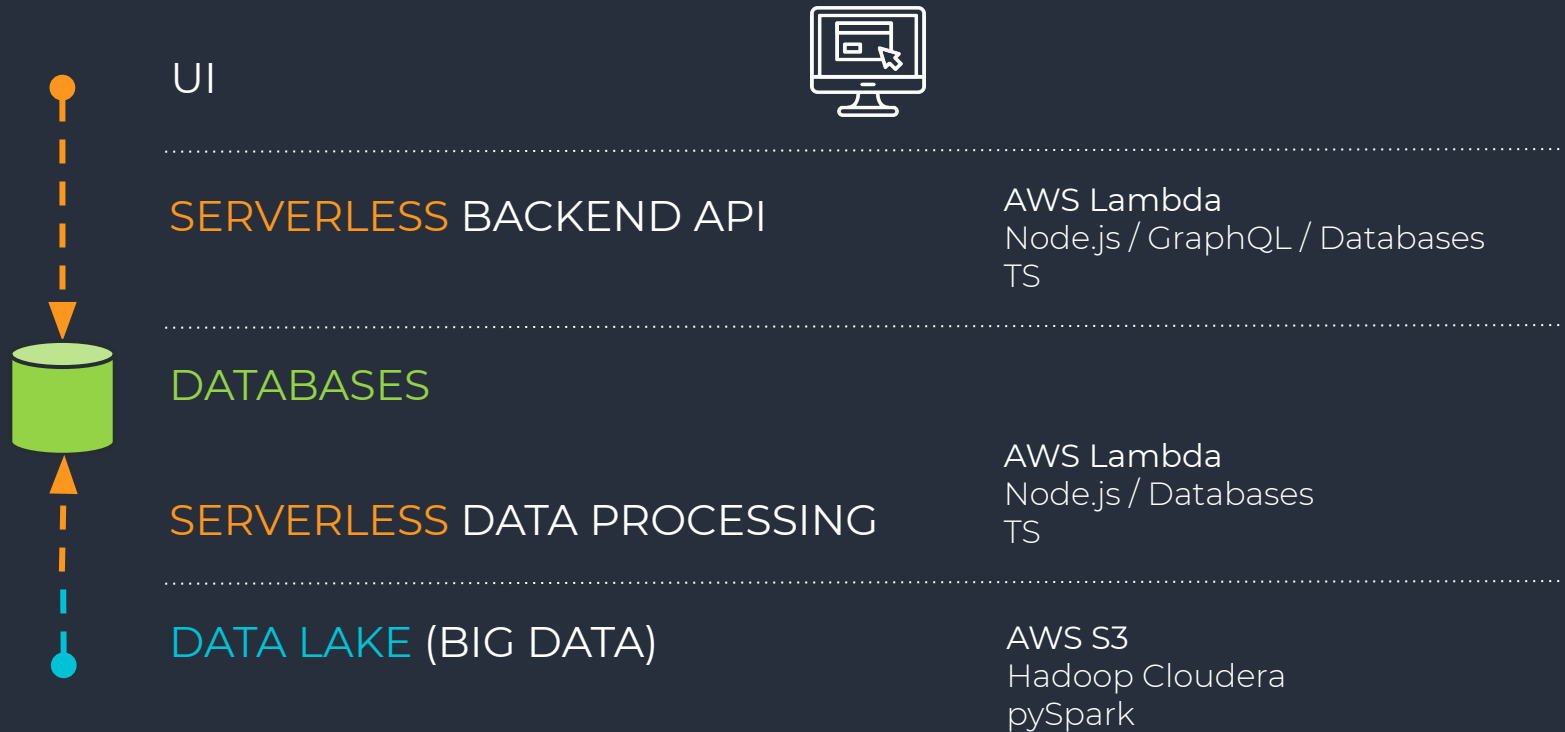
- Hundreds GB RAW data and 60M new messages daily
- 8M UNIQUE ITEMS over the world



TECHNOLOGIES



TECHNOLOGIES



SERVERLESS FUNCTION

```
export const handler = async (event) => {  
  const data = event.Records[0].body;  
  
  // - TRANSFORM data  
  // - WRITE to DB or  
  // - PUT TO QUEUE/STREAM/TOPIC  
  
  return 'success';  
};
```




- NO INFRASTRUCTURE TO MANAGE
- TRIGGERED BY EVENTS
- HIGH SCALABLE
- STATELESS
- COST-EFFECTIVE
- CHOOSE YOUR CODE LANGUAGE

NODE.JS POPULARITY

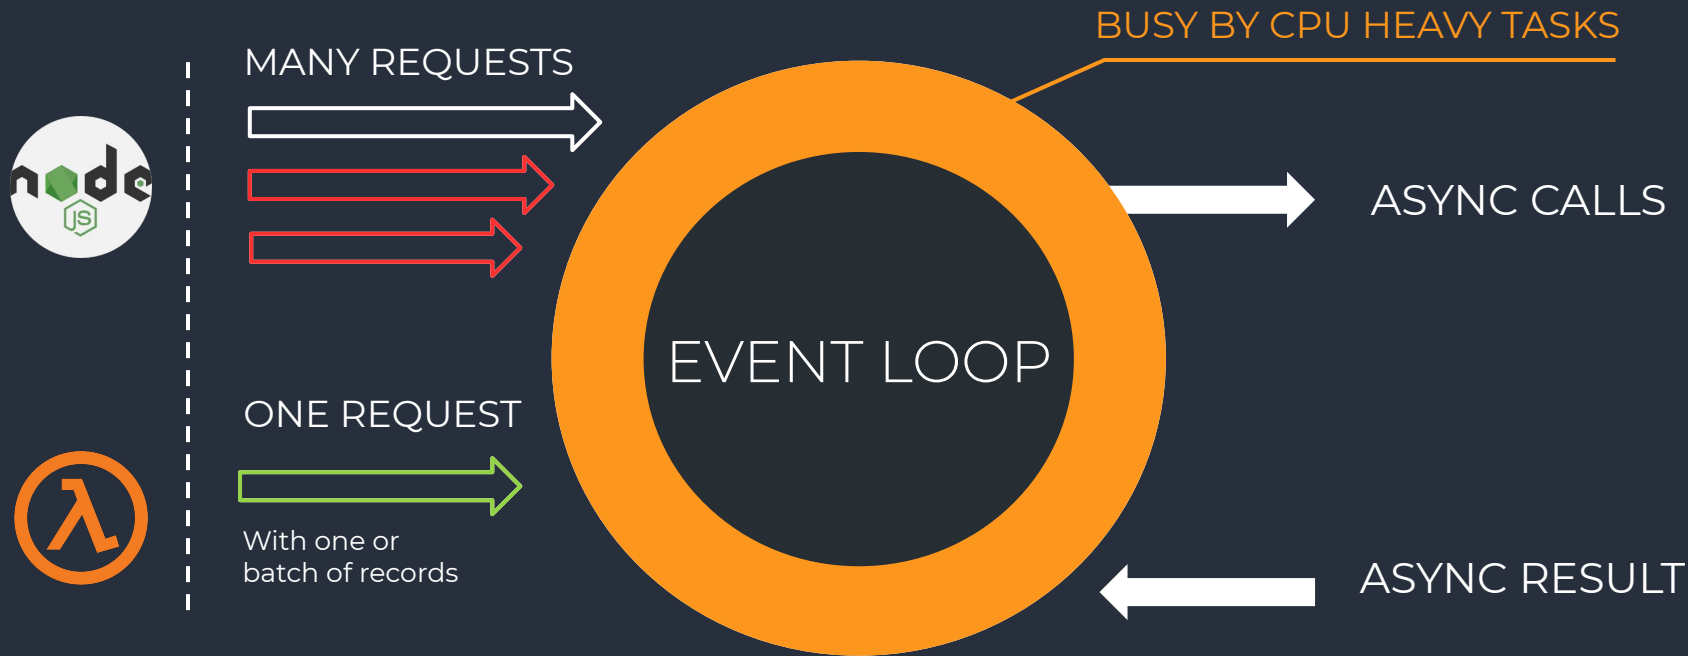
2014 AWS LAMBDA WITH NODE.JS

- PURE ASYNCHRONOUS
- MINIMALISTIC CORE
- FAST STARTUP WITH HIGH PERFORMANCE



2018 SERVERLESS.COM

NODE.JS FUNCTION VS SERVER



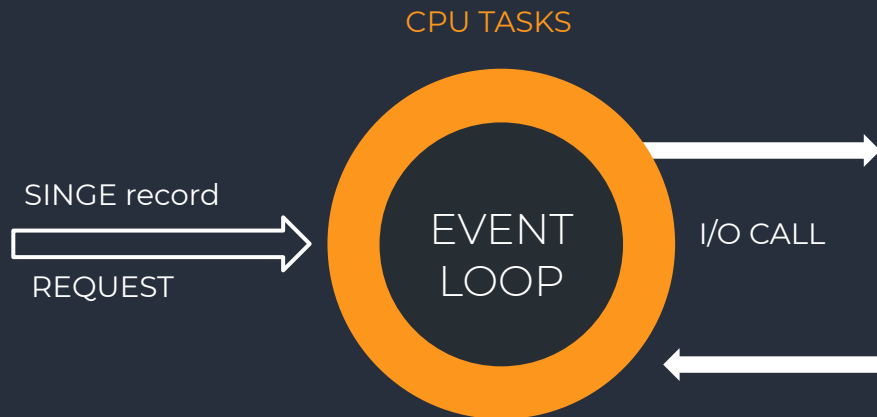
I/O – PERFORMANCE BOTTLENECK



SUMMARY


CPU HEAVY TASKS

Get request with single record



I/O INTENSIVE TASKS

Get request with batch of records (like a server)




SERVERLESS COMPUTE SERVICE

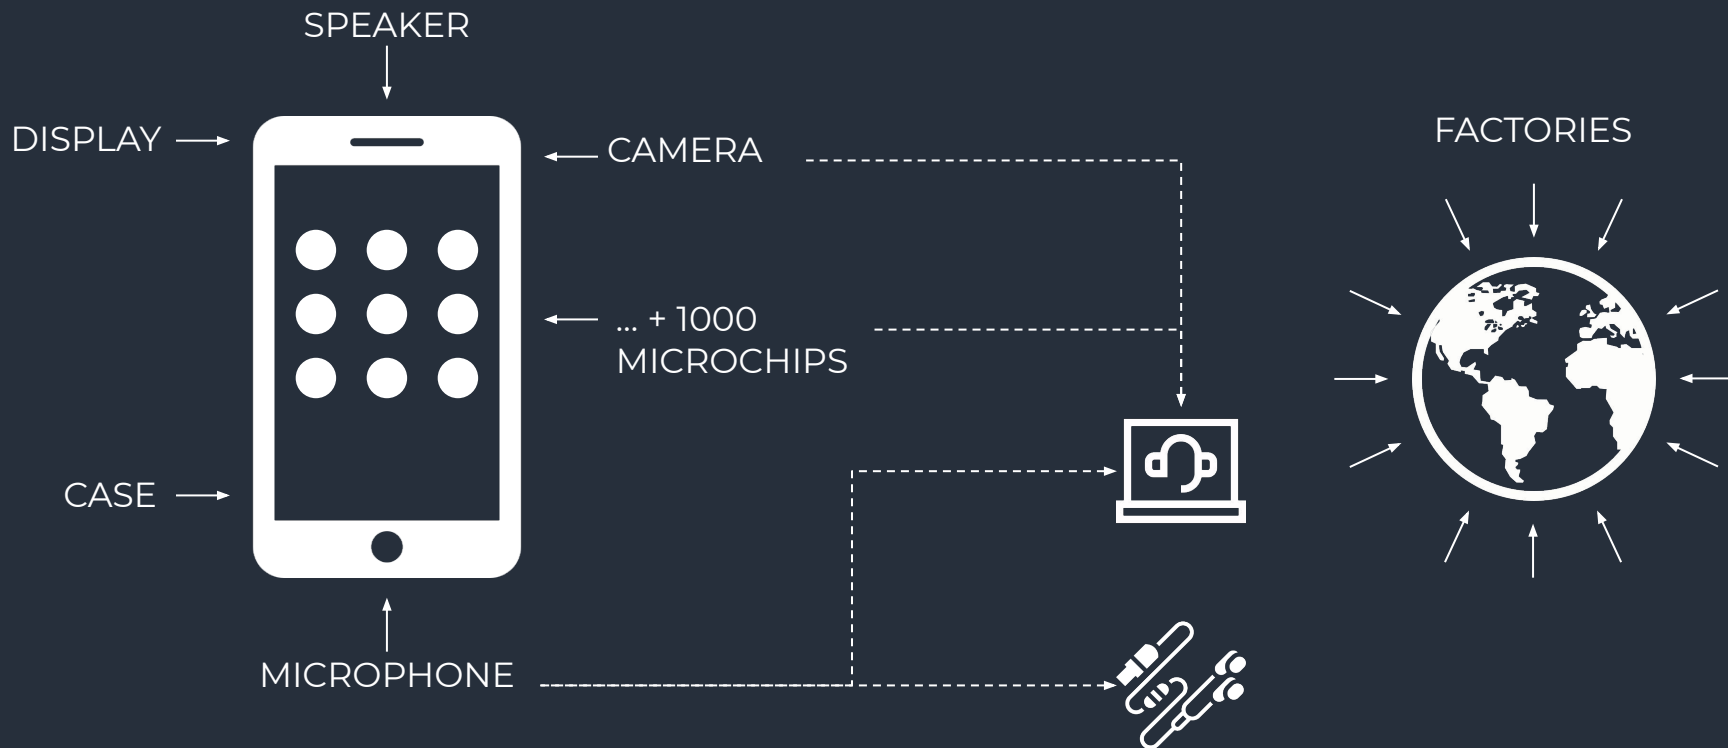


THIS IS MY ARCHITECTURE – 177 VIDEOS
<https://aws.amazon.com/ru/this-is-my-architecture/>


Symbol in the presentation.



EXAMPLE



BATCH PROCESSING

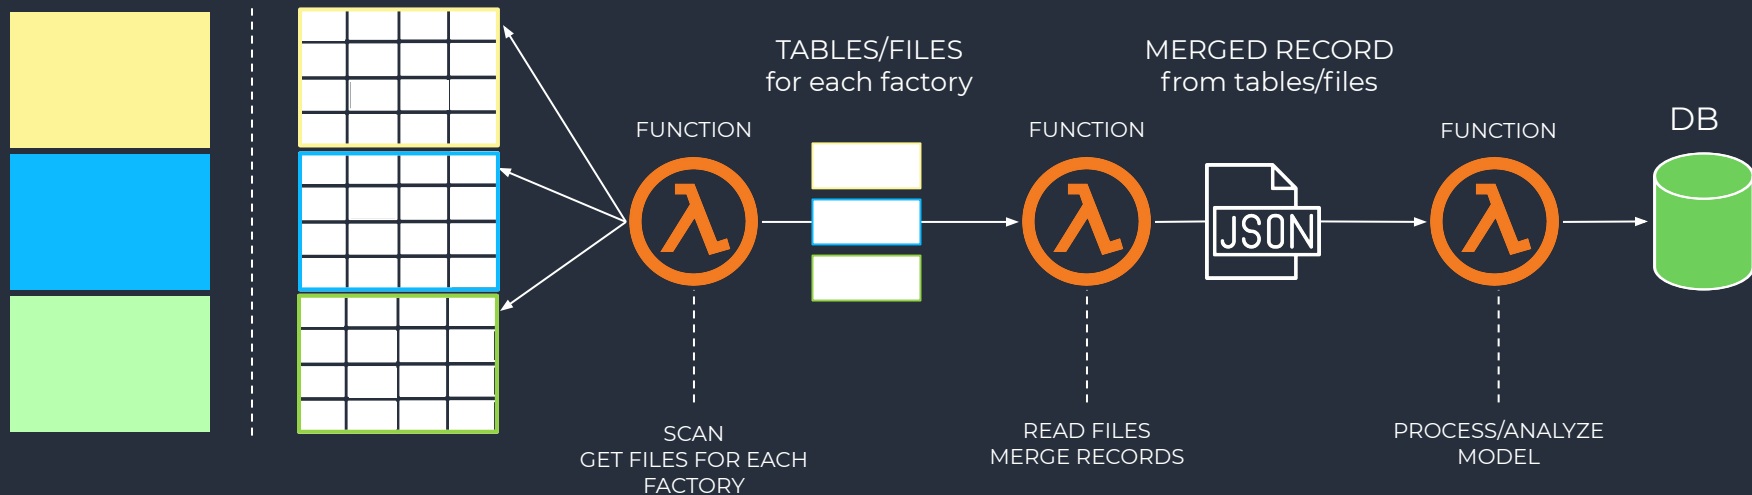


BATCH DATA PROCESSING

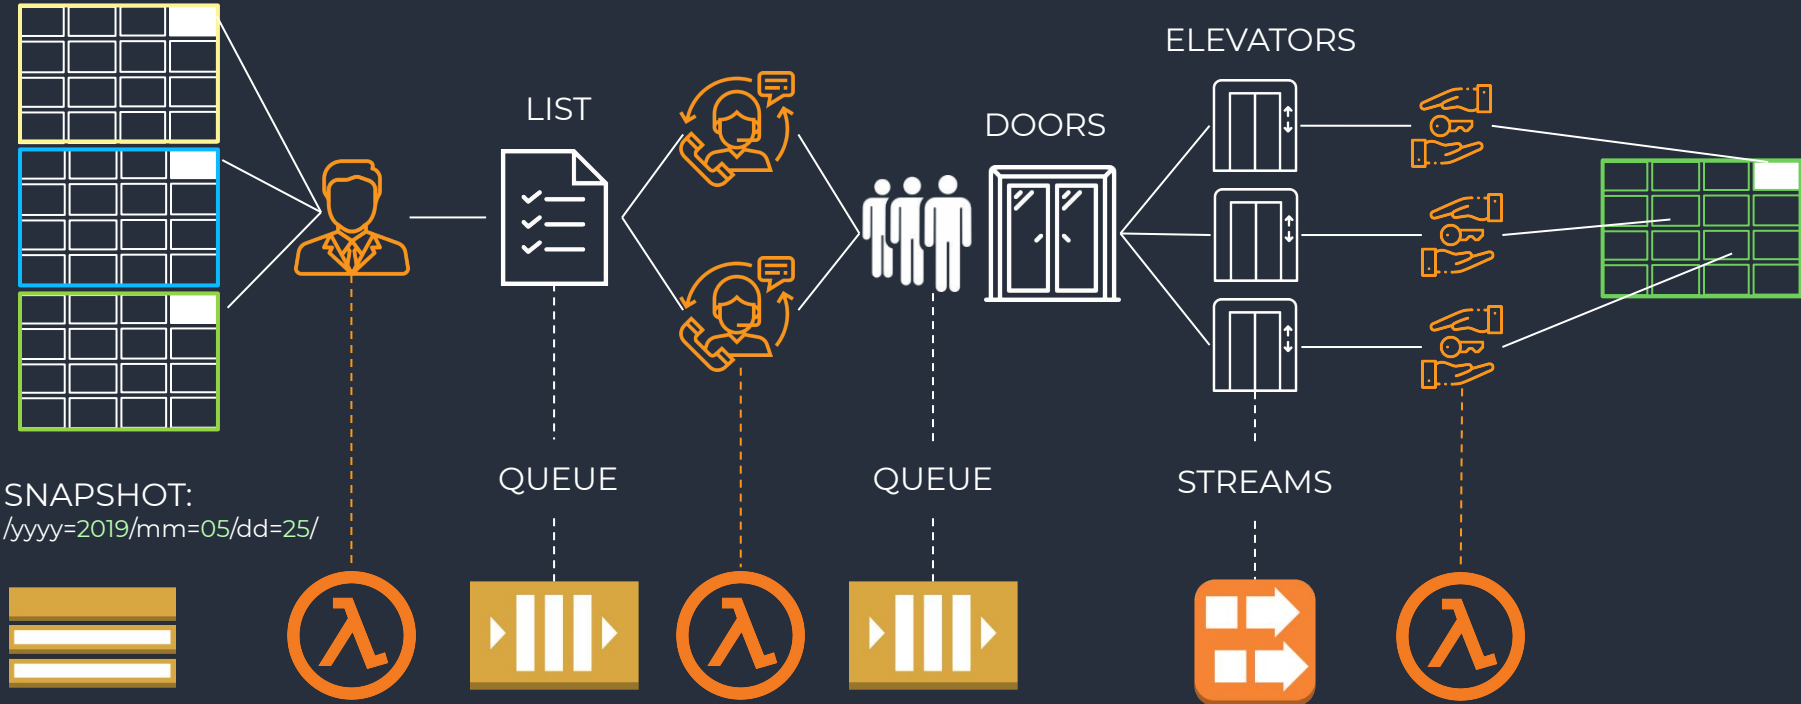
SNAPSHOT:

/yyy=2019/mm=05/dd=25/

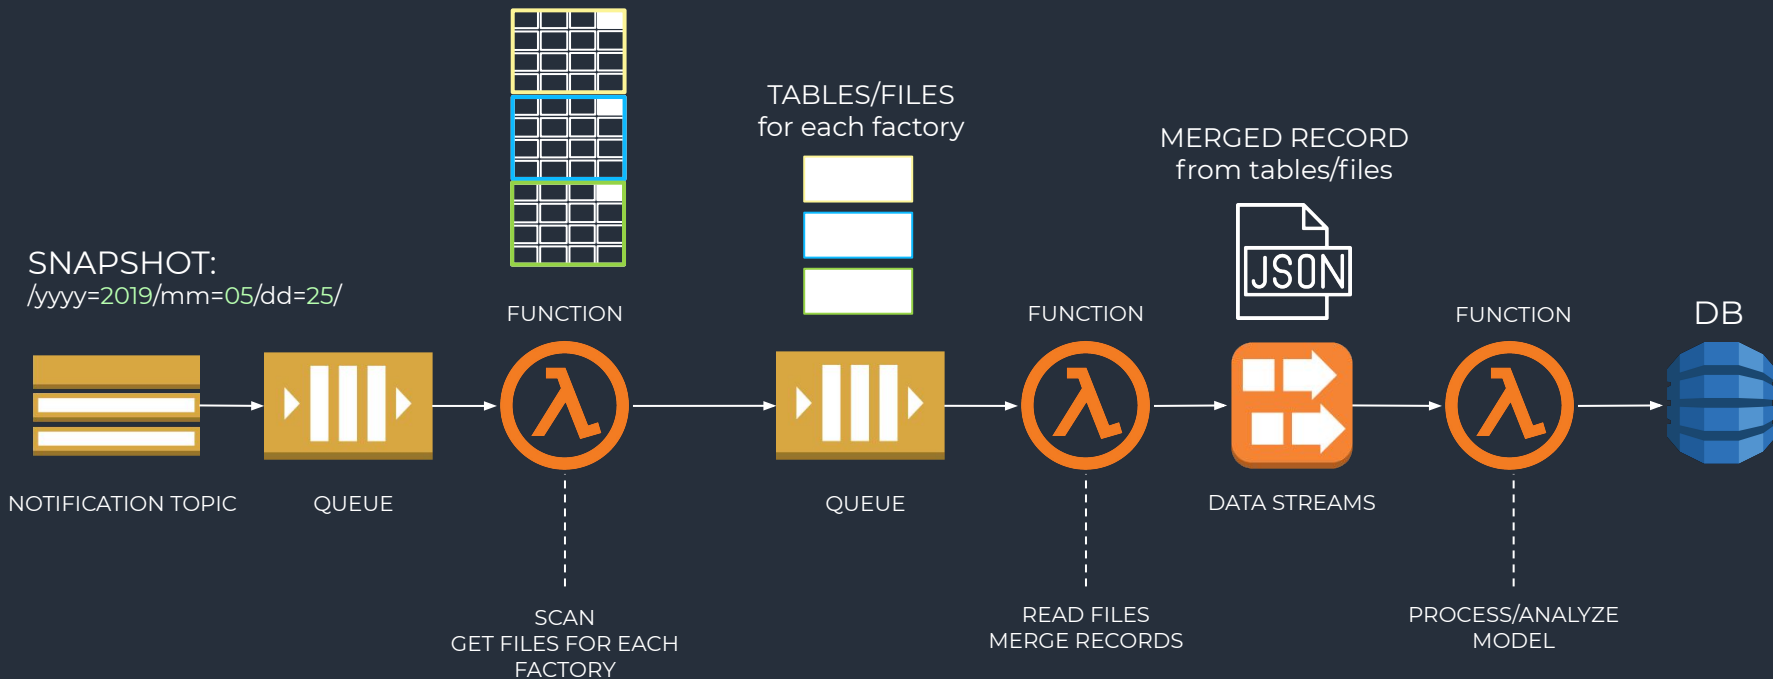
→ DONE!




RELOCATION OF PEOPLE TO A NEW BUILDING



BATCH PROCESSING ARCHITECTURE



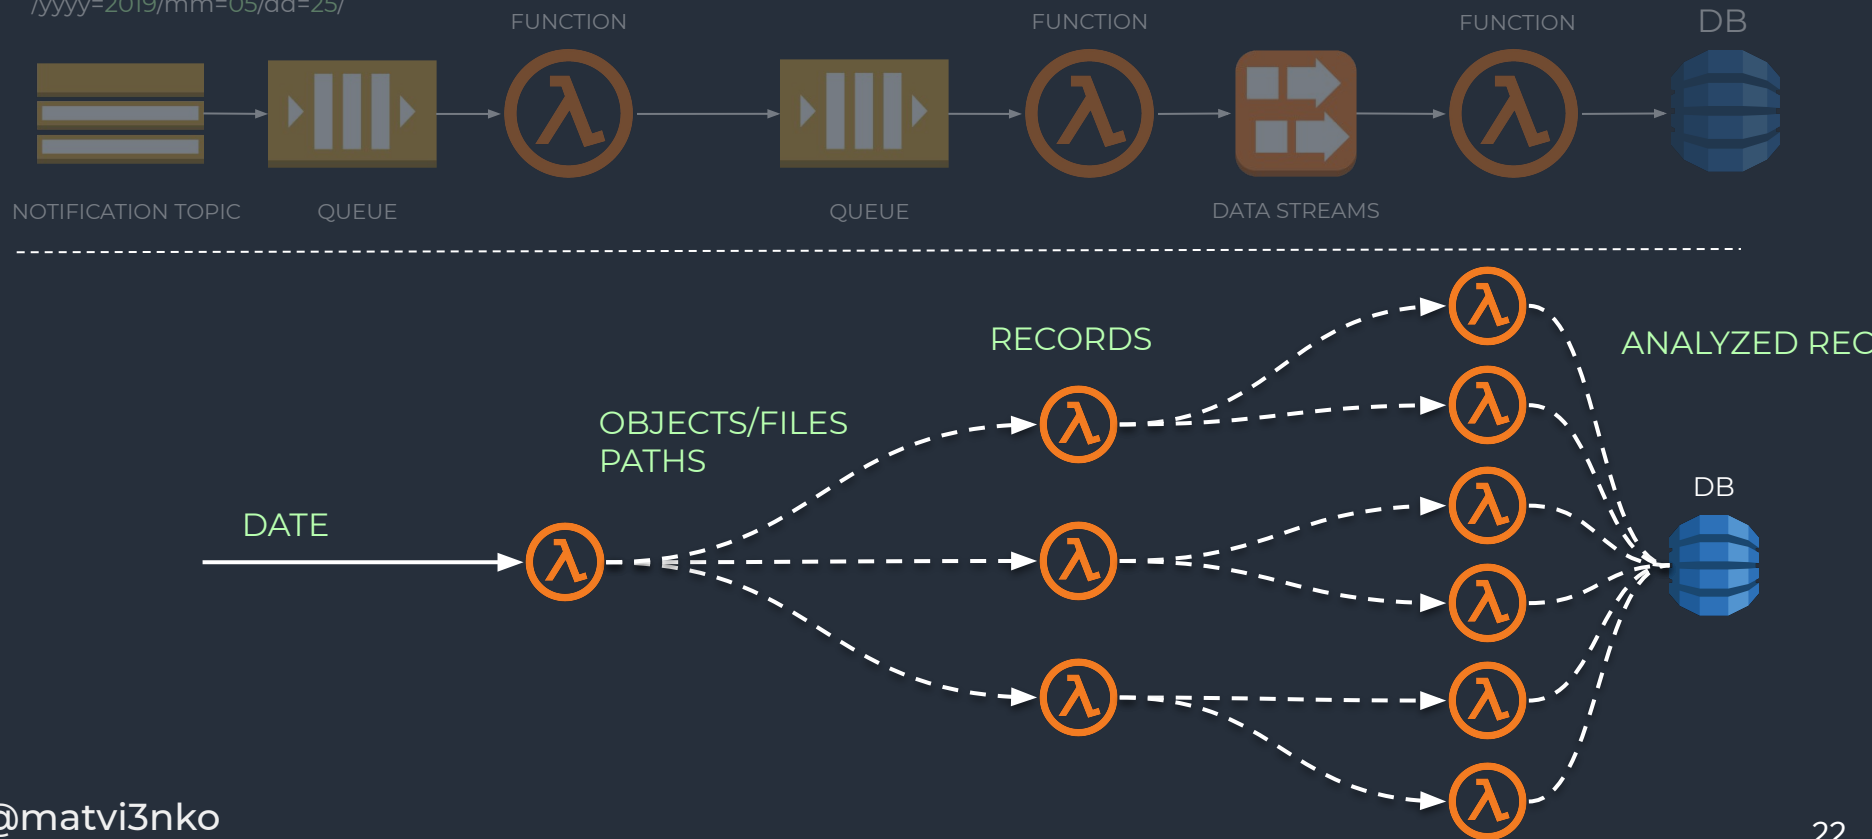
REAL-TIME PROCESSING ARCHITECTURE



SCALABILITY

SNAPSHOT:

/yyyy=2019/mm=05/dd=25/



CONCLUSION

1. FROM BIG DATA TO A LARGE NUMBER OF MESSAGES
2. THE MORE MESSAGES THE FUNCTION ACCEPTS, THE MORE IT NEEDS TO BE PARALLELIZED
3. USE THE QUEUE FOR MESSAGES, AND DATA STREAMS TO TRANSFER MODELS / LARGE COLLECTION
4. INCREASE THE NUMBER OF STREAM SHARDS. 1 SHARD = 1 LAMBDA FUNCTION
5. PREPARE TO STREAMING / REAL-TIME PROCESSING

PROGRESS



SOLUTION STRUCTURE AND FUNCTION

BASE ARCHITECTURE DESIGN

SERVERLESS PROJECT STRUCTURE

/transform

– serverless.yml

– handler.ts

/analyze

– serverless.yml

– handler.ts

/node_modules

serverless.yml

package.json

```
import AWS from 'as-sdk';
```

```
const s3Client = new AWS.S3({region});
```

```
export const handler = (event) => {
```

```
  const [message] = event.Records;
```

```
  return new Promise((resolve, reject) => {
```

```
    this.s3Client.selectObjectContent({ Key: message.path }, (err, data) => {
```

```
      if (err) {
```

```
        reject(err);
```

```
      }
```

```
      resolve(data);
```

```
    });
```

```
  };
```

SERVERLESS PROJECT STRUCTURE

/transform

- serverless.yml
- handler.ts

/analyze

- serverless.yml
- handler.ts

/node_modules

serverless.yml
package.json

DISADVANTAGES

1. NODE_MODULES contains dependencies of all functions
Have to control and split them in SERVERLESS.YML
2. Lack of function isolation
3. Lack of independent install / build / test
4. Becomes monolith project

MONOREPO SERVERLESS PROJECTS STRUCTURE

/lib
 /node_modules
 /errors
 /factories
 /models
 /providers

– package.json

/transform(er)

 /node_modules

– functionA1.ts

– package.json

– serverless.yml

/analyze(r)

 /node_modules

– functionA2.ts

– package.json

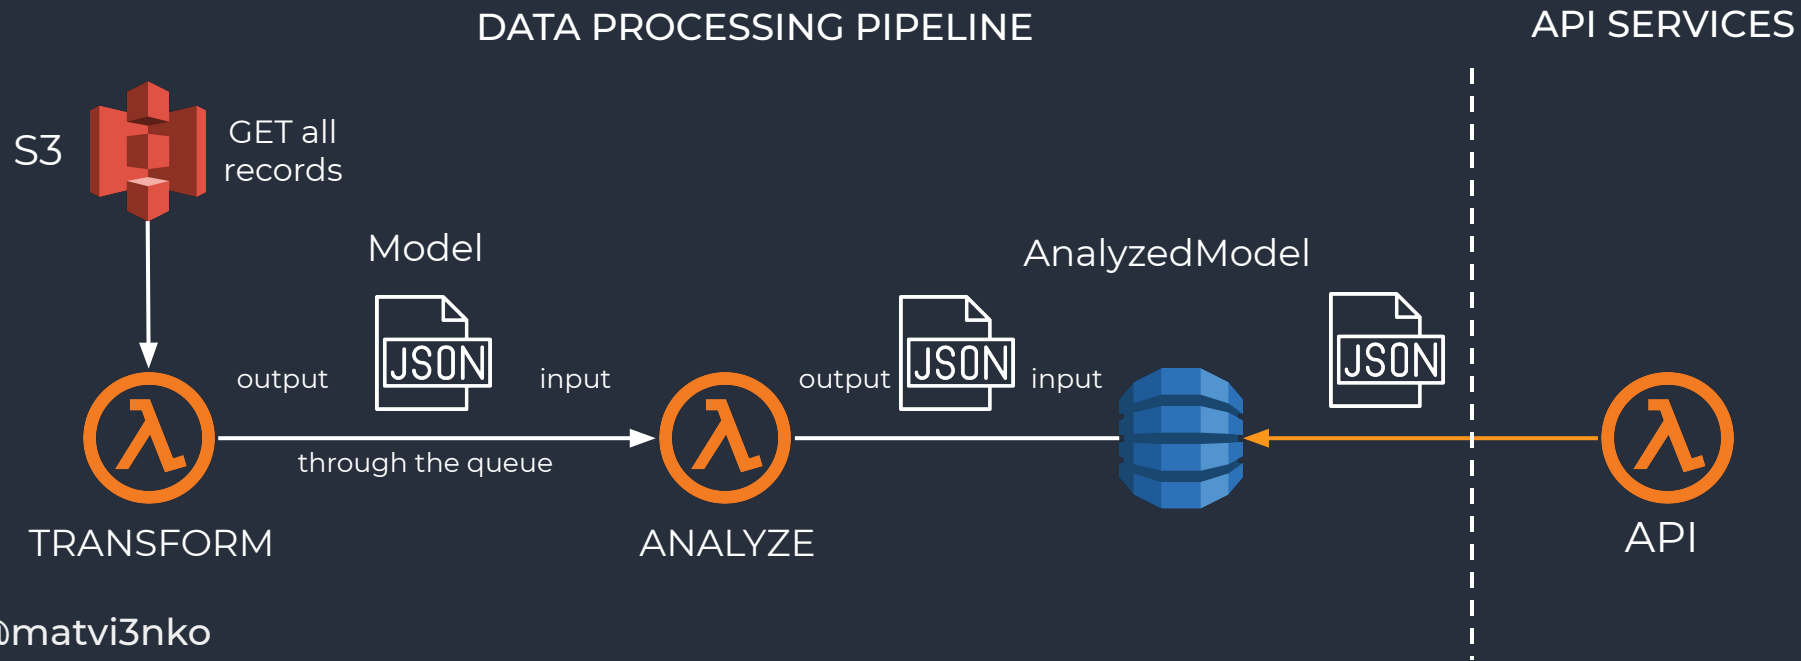
– serverless.yml

ADVANTAGES

1. LIB contains all common infrastructure, domain logic and cloud provider's SDK
2. Functions became isolated projects with flexible splitting and contains only business logic
3. LIB and PROJECTS versioning
4. NPM resolves NODE_MODULES dependencies automatically
5. Independent install / build / test / deploy / troubleshooting

MODELS REUSABILITY

```
import { Model } from '@holyls/models'
```



INITIALIZE ONLY ONCE

FIRST EXECUTION
ON COLD START

```
import { createCsvS3Provider } from '@holyljs/lib';  
const service = new Service(createCsvS3Provider());
```

ALL NEXT
ON WARM START

```
export const handler = async (event) => {  
  // Logic  
};
```

ENCAPSULATE IN LIB
BOILERPLATE CODE
USE FACTORIES / IoC

USE STREAMING PROCESSING

USE RXJS PIPES TO

EXTRACT
↓
TRANSFORM
↓
SEND

```
import { createCsvS3Provider, createQueueProvider } from '@holyls/lib';
const service = new Service(createCsvS3Provider());

export const handler = async (event) => {
  const source = service.getObject(event)
    .pipe(flatMap(service.transform))
    .pipe(bufferCount(10))
    .pipe(flatMap(message => createQueueProvider().putBatch(message)));

  return new Promise((resolve, reject) => {
    source.subscribe(() => { /* handle */ },
      err => {
        err instanceof InfrastructureError && reject(err);
        err instanceof DomainError && reject(err);
      },
      resolve);
  });
};
```

HANDLE ERRORS AT A HIGH LEVEL

1. CHECK ERROR TYPE

2. DECIDE WHAT TO DO

3. DECIDE WHERE TO SEND

```
import { createCsvS3Provider, createQueueProvider } from '@holyls/lib';
const service = new Service(createCsvS3Provider());

export const handler = async (event) => {
  const source = service.getObject(event)
    .pipe(flatMap(service.transform))
    .pipe(bufferCount(10))
    .pipe(flatMap(message => createQueueProvider().putBatch(message)));

  return new Promise((resolve, reject) => {
    source.subscribe(() => { /* handle */},
      err => {
        err instanceof InfrastructureError && reject(err);
        err instanceof DomainError && reject(err);
      },
      resolve);
  });
};
```

4 COMPONENTS OF THE FUNCTION

1. INCOMING DATA
2. TYPE OF TRANSFORMATION
3. DESTINATION
4. MAIN ERROR HANDLER

PROGRESS




ERROR HANDLING

SOLUTION STRUCTURE AND FUNCTION

BASE ARCHITECTURE DESIGN

TRACK LOG MESSAGES THAT ARE HIDDEN ERRORS

1. Request XX-YY: "Process exited before completing request"
2. Function completed on its timeout (up to limit)



HOW TO HANDLE

1. FUNCTION HANGS

Don't do that: `context.callbackWaitsForEmptyEventLoop = false;`

or close Sequelize connection to fix that

Use callback cb(), rewrite to async/await (Promises)

2. FUNCTION DOES NOT PERFORM PART OF THE LOGIC

You added async or return value.

Find missed await / return Promise



WAIT FOR RESPONSES FROM THE SERVICES

// CODE

SEND(MESSAGE);

// CODE

// CODE

RETURN SEND(MESSAGE);

// CODE

// CODE

AWAIT SEND(MESSAGE);

// CODE

SEND



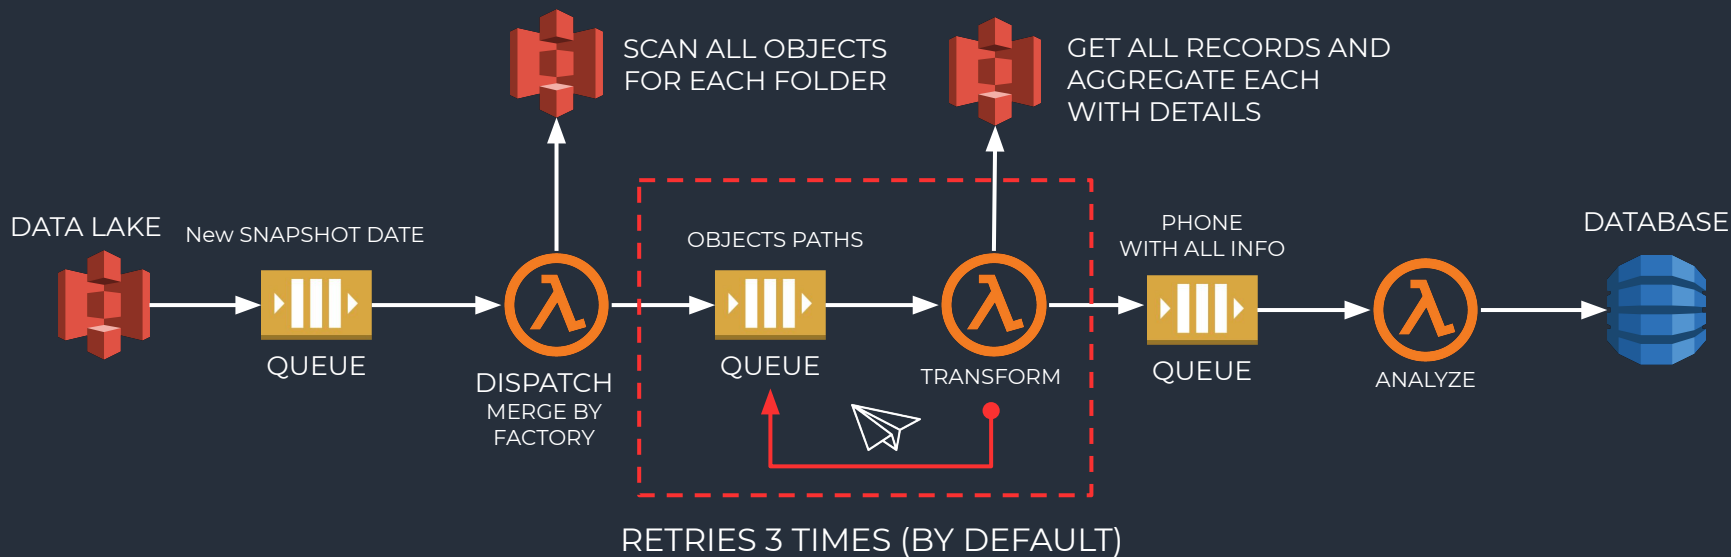
SAVE 1\$ → LOSE MILLIONS \$



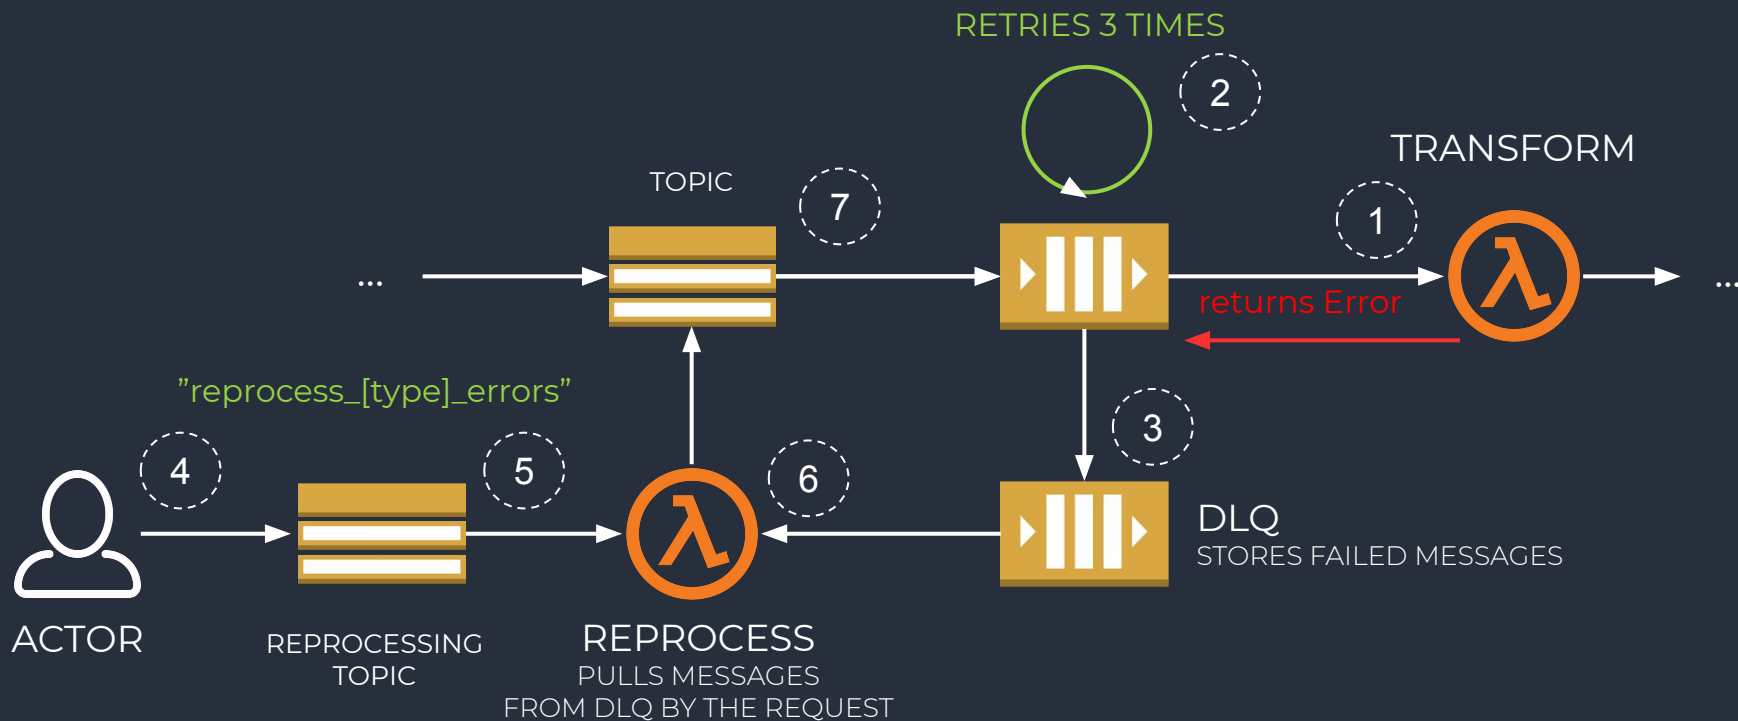
ERROR!



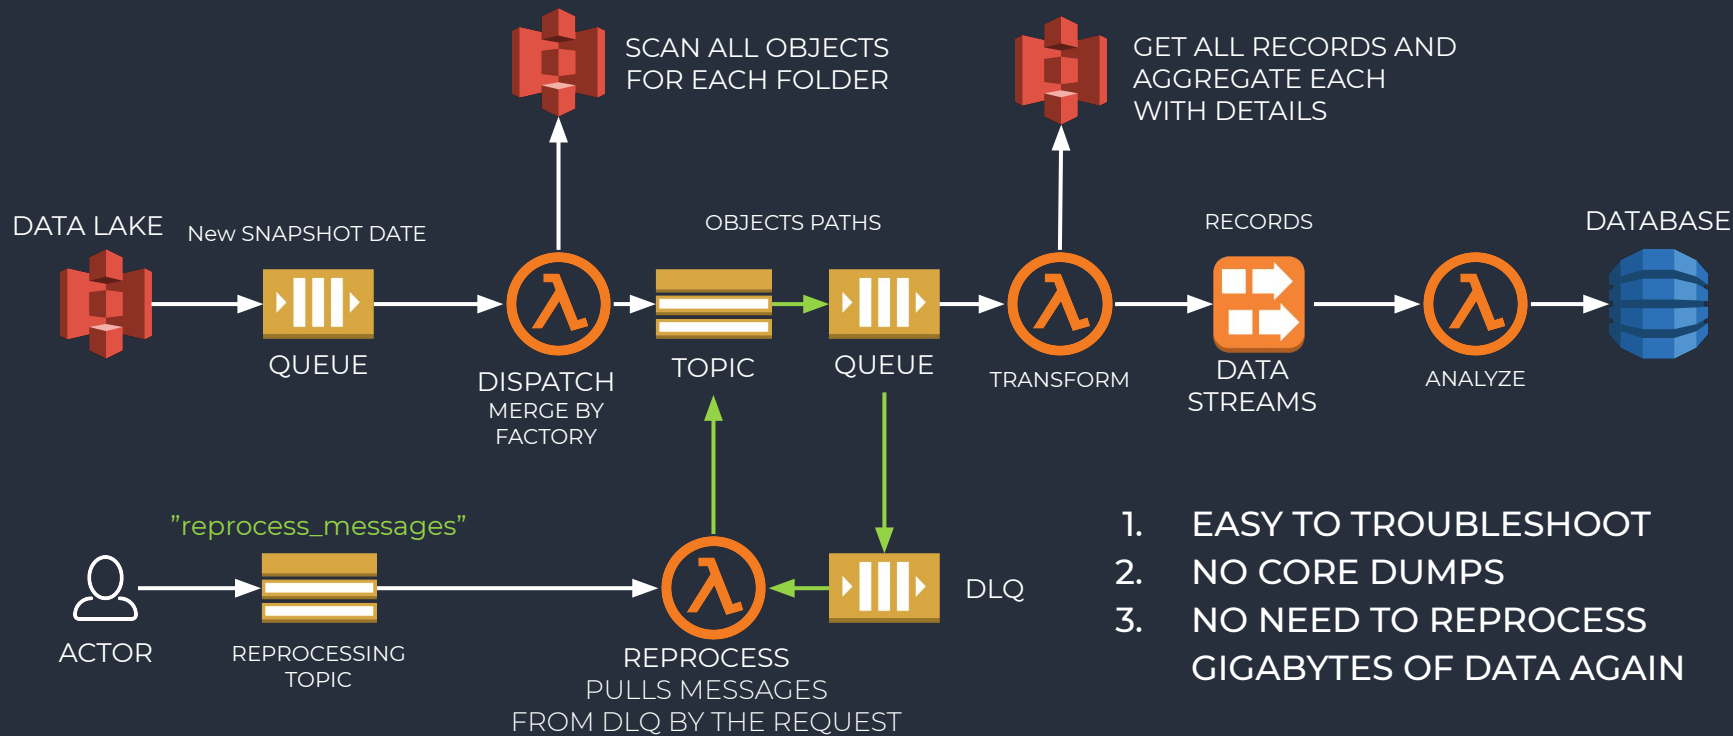
RETRY STRATEGY



DEAD LETTER QUEUE

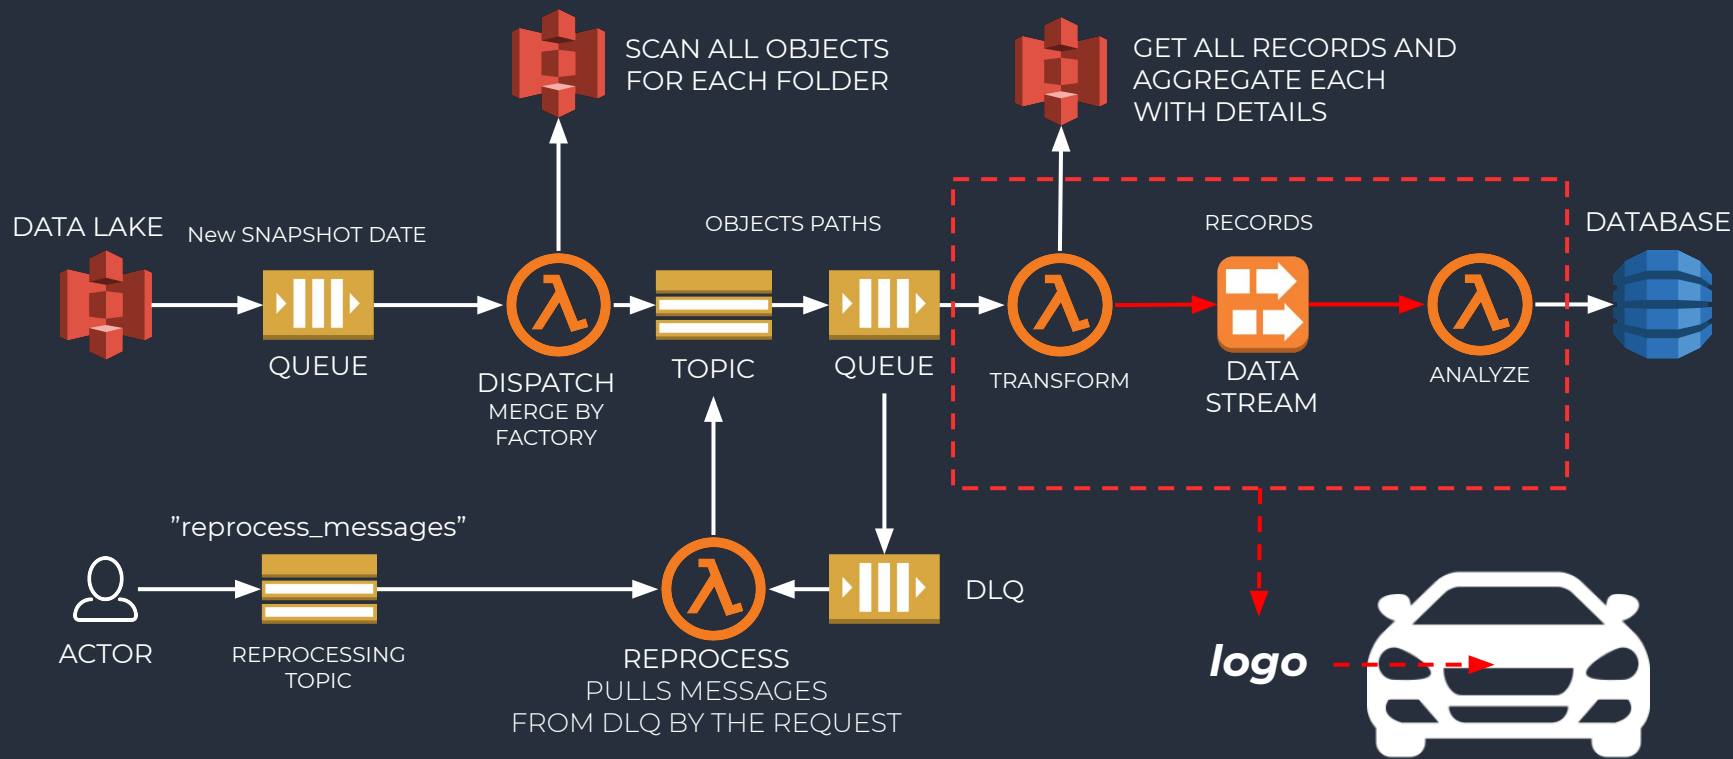


DLQ FOR THE QUEUE



1. EASY TO TROUBLESHOOT
2. NO CORE DUMPS
3. NO NEED TO REPROCESS GIGABYTES OF DATA AGAIN

KINESIS ERROR HANDLING



ERROR HANDLING STRATEGY

1. PUT TO STREAM (TRANSFORM)

- a. Handle partially successful response

```
FailedRecordCount : Number
```



2. READ FROM (ANALYZE)

- a. Reduce retry times


```
customBackoff: (retryCount) => {...}
```

- b. Use Dead Letter Queue

- c. Handle duplicate Records



UP TO 7 days each 100 ms



EXACTLY ONCE PROCESSING

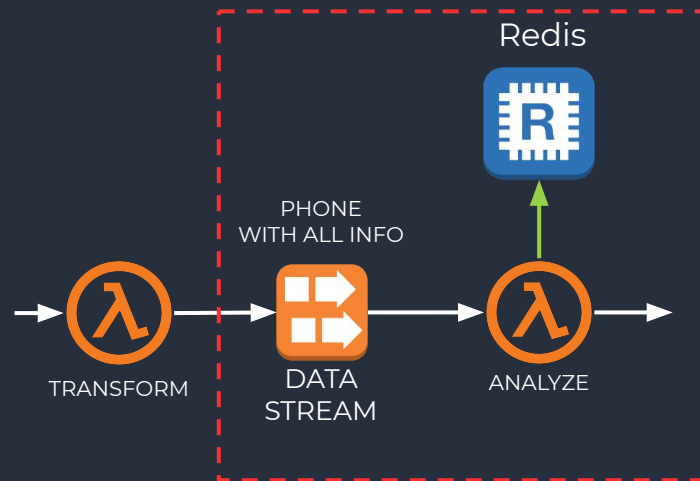
1. REDUCE THE RISK OF FAILURE

```
const response = await putToStream(record);  
// do something with response -> RISK
```

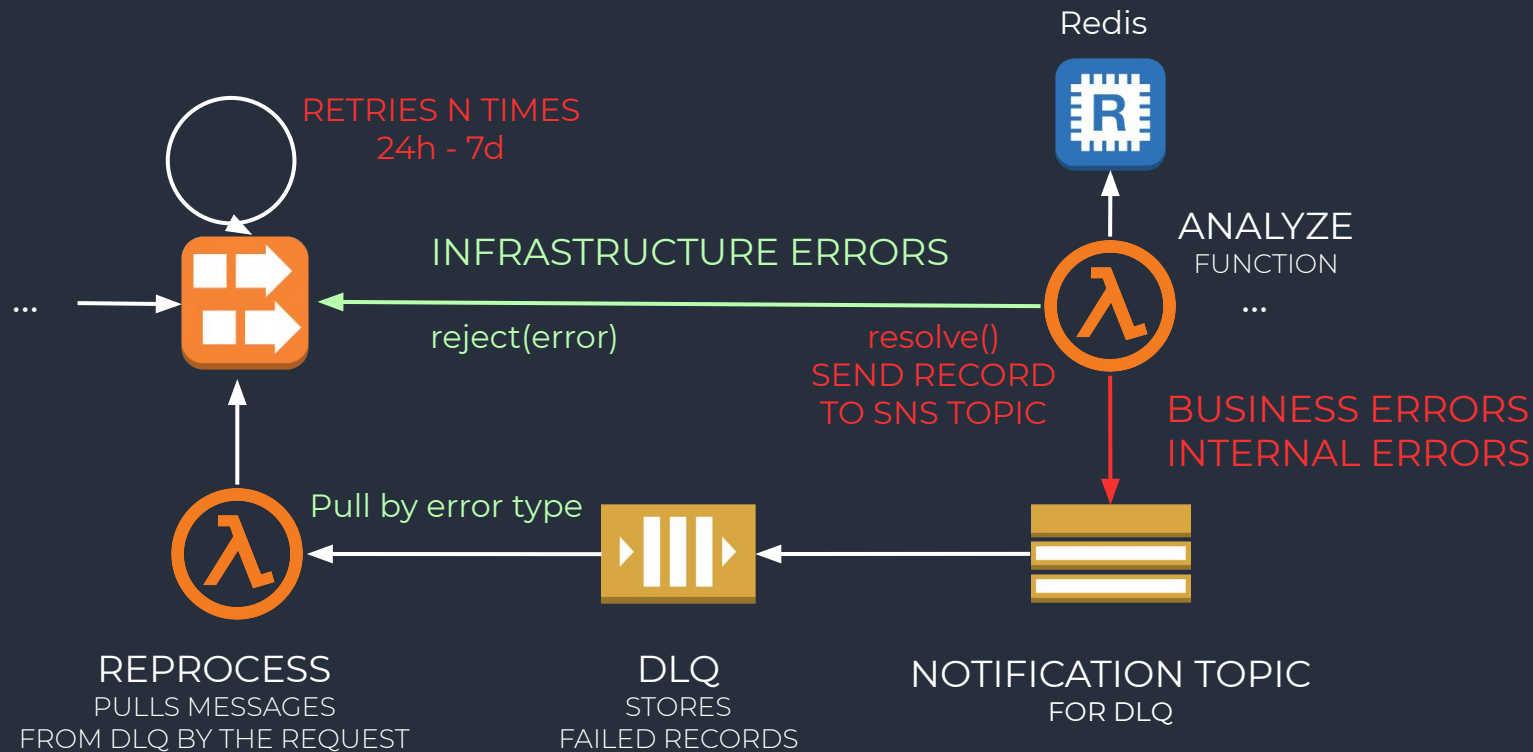
2. USE REDIS CACHE TO STORE KEYS OF RECORDS

key: [date]-[shard-id]-[sequence-number]

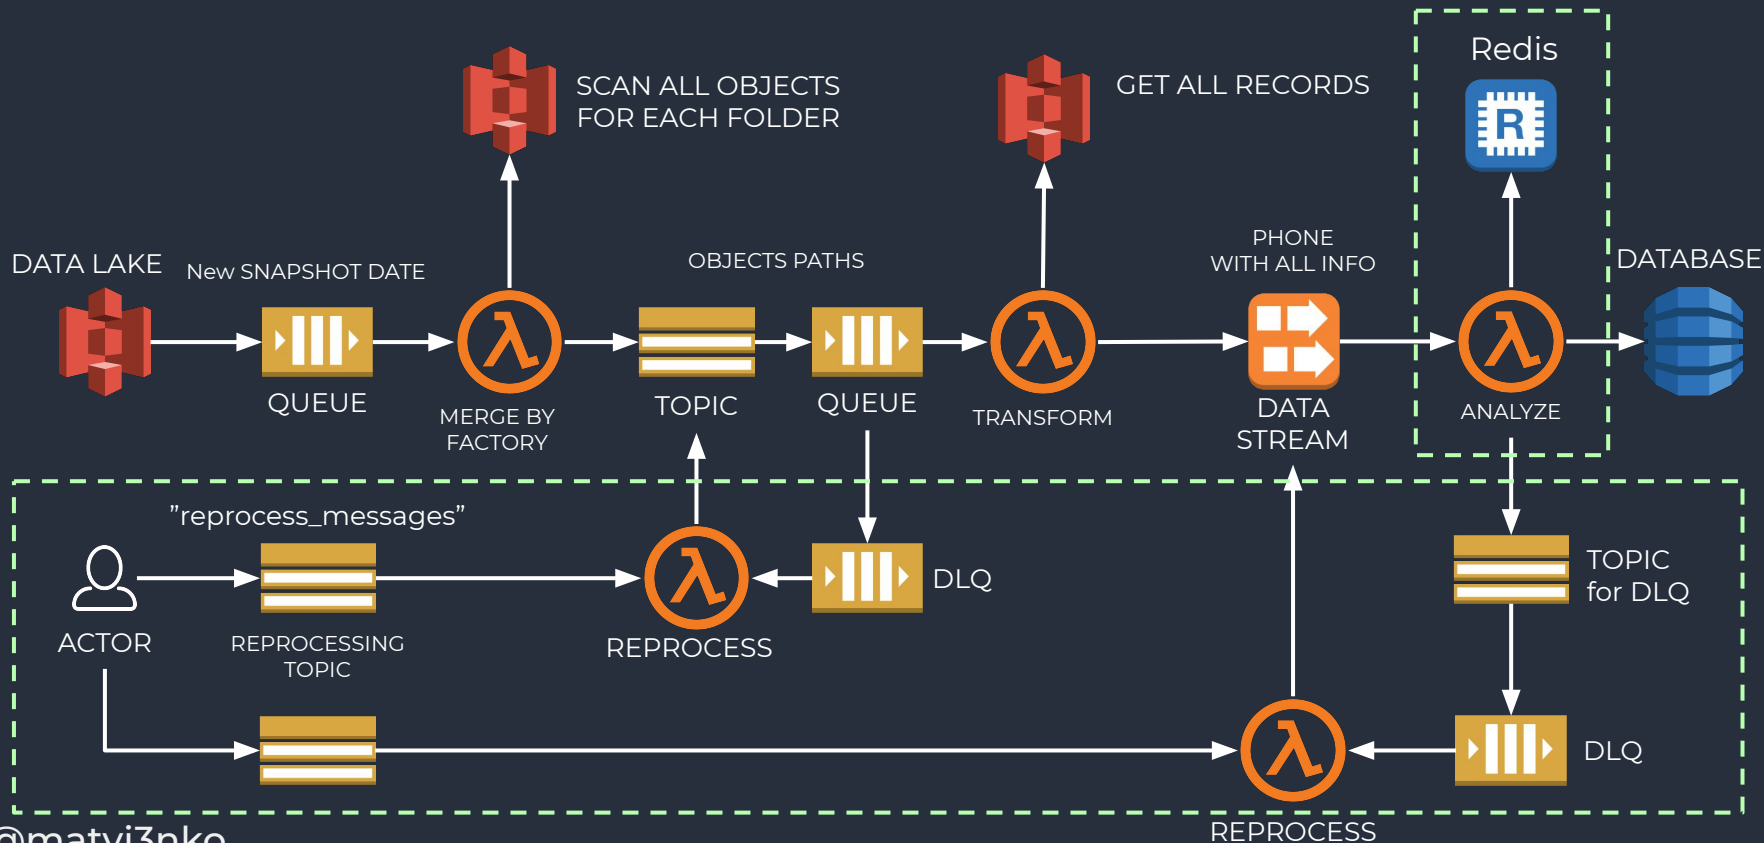
3. FILTER DUPLICATES



KINESIS ERROR HANDLING



ARCHITECTURE WITH REPROCESSING BLOCK



CONCLUSION

1. CREATE CUSTOM TYPES OF ERRORS
2. REPROCESS ONLY FAILED MESSAGE AND NOT GIGA/TERABYTES OF DATA
3. DON'T LOSE MESSAGES
4. USE DLQ WITH FILTERING
5. PROCESS EXACTLY ONCE

PROGRESS



READ AND PROCESS DATA

ERROR HANDLING


SOLUTION STRUCTURE AND FUNCTION

BASE ARCHITECTURE DESIGN

AWS S3: SIMPLE STORAGE SERVICE

1. S3 SELECT REQUEST

```
const params = {  
  Bucket: 'bucket_name',  
  Key: 'key_name',  
  ExpressionType: 'SQL',  
  Expression: `  
    SELECT s.name, s.year  
    FROM S3Object s  
    WHERE s.name = 'HolyJS'  
  `,  
},  
OutputSerialization: {  
  JSON: { RecordDelimiter: '\n' },  
};
```



JSON | CSV | Parquet



2. QUERY RESULT

id	name	year	details

UP TO 4x FASTER

DATA FORMATS



RAW OR COMPRESSED?



80MB
PARQUET
BINARY



1.4GB
CSV
TEXT. ONLY VALUES

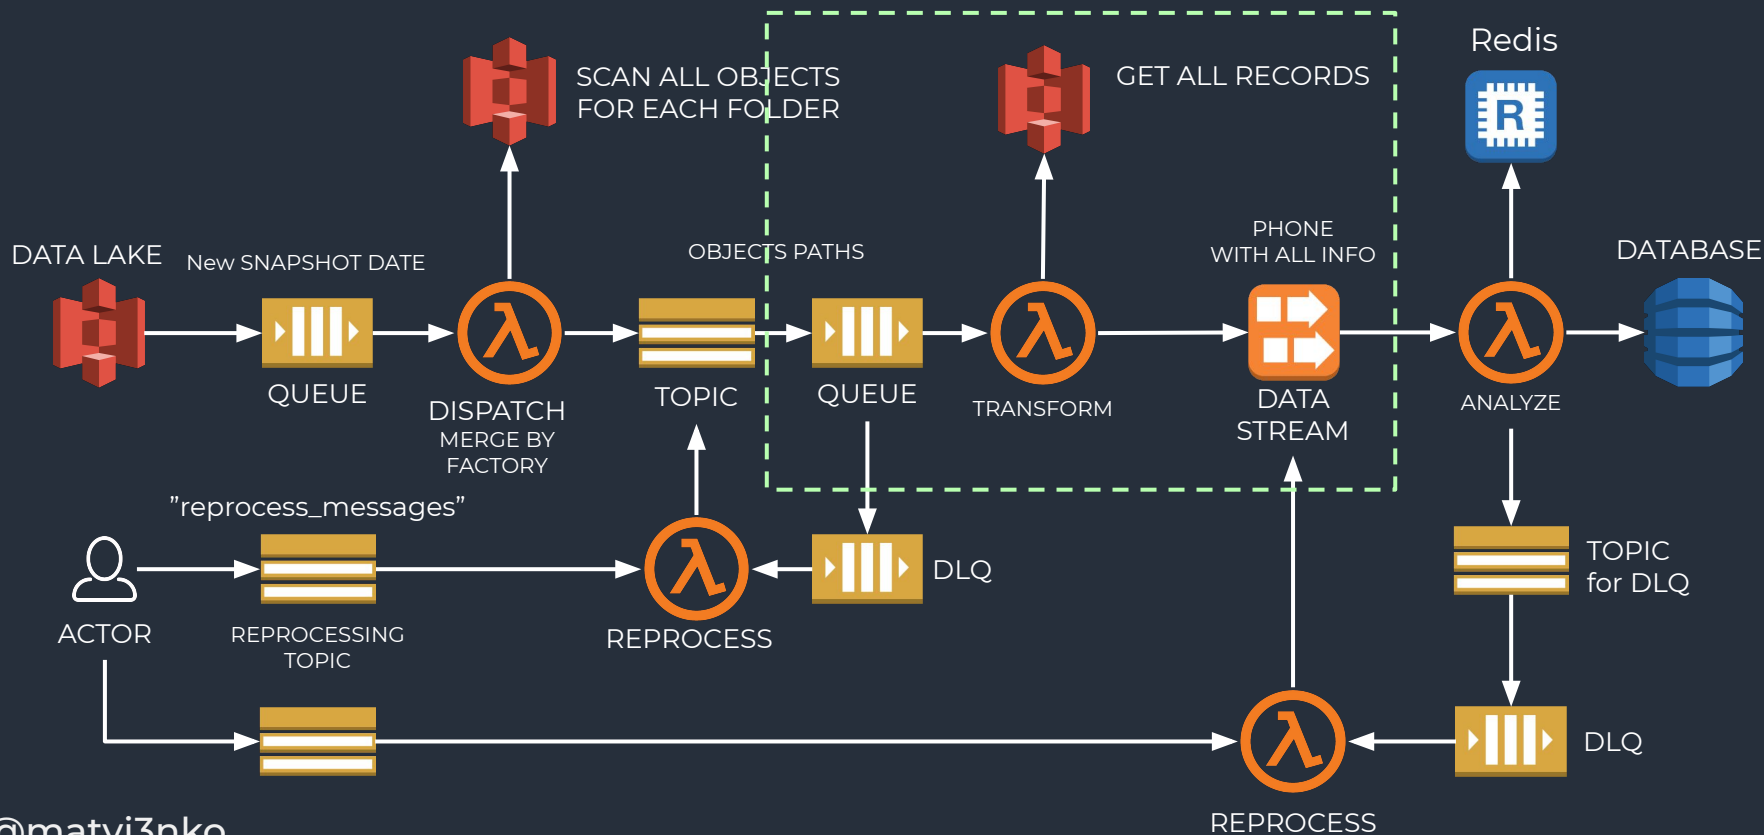
~X20 bigger size



6.5GB
JSON
TEXT. FIELDS + VALUES

~X80 bigger size
75% - fields names

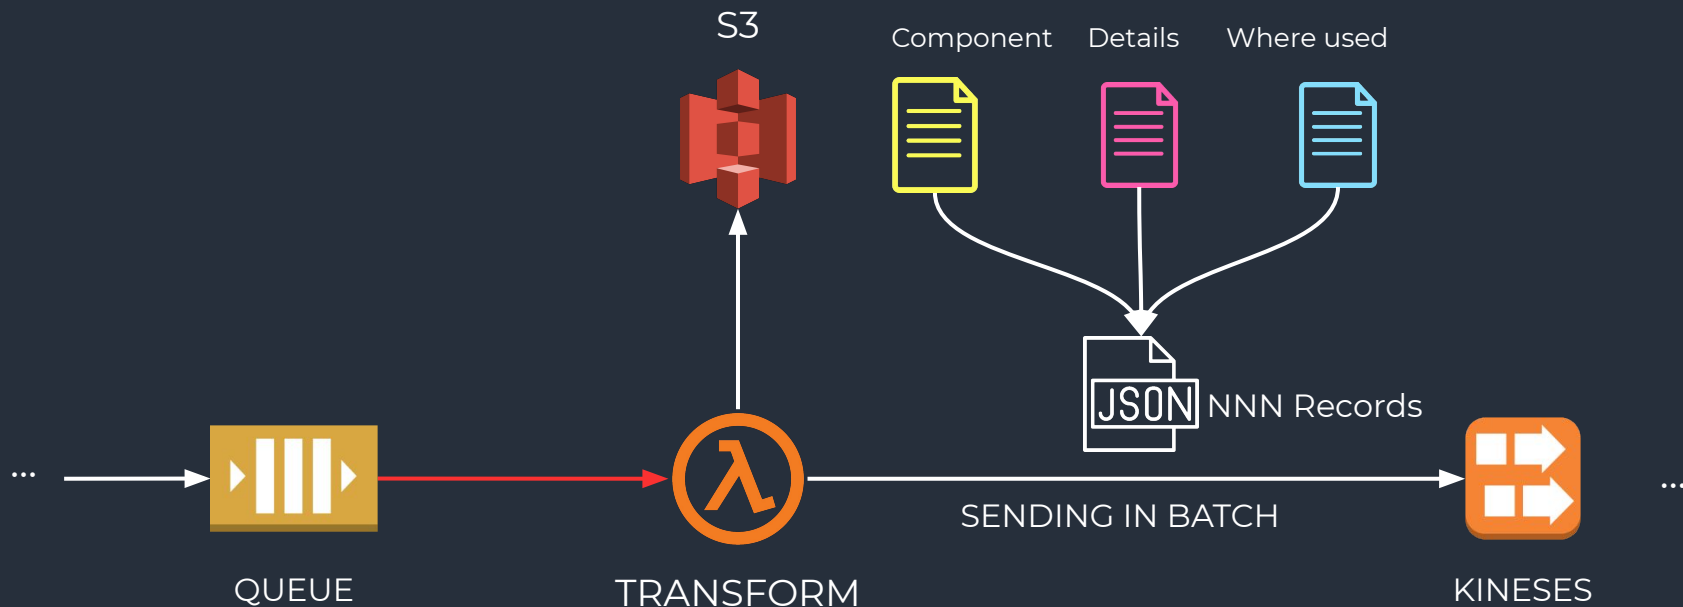
EXTRACTING BIG OBJECTS/FILES



EXTRACT DATA


S3 OBJECT PATH:

s3/buckets/bucket-name/**entity**/yyyy=**2019**/mm=**05**/dd=**25**/**partition-by-category**/key.parquet




DATA TRANSFORMATIONS

Component Details Where used




ELEMENT-WISE

Component Details Where used



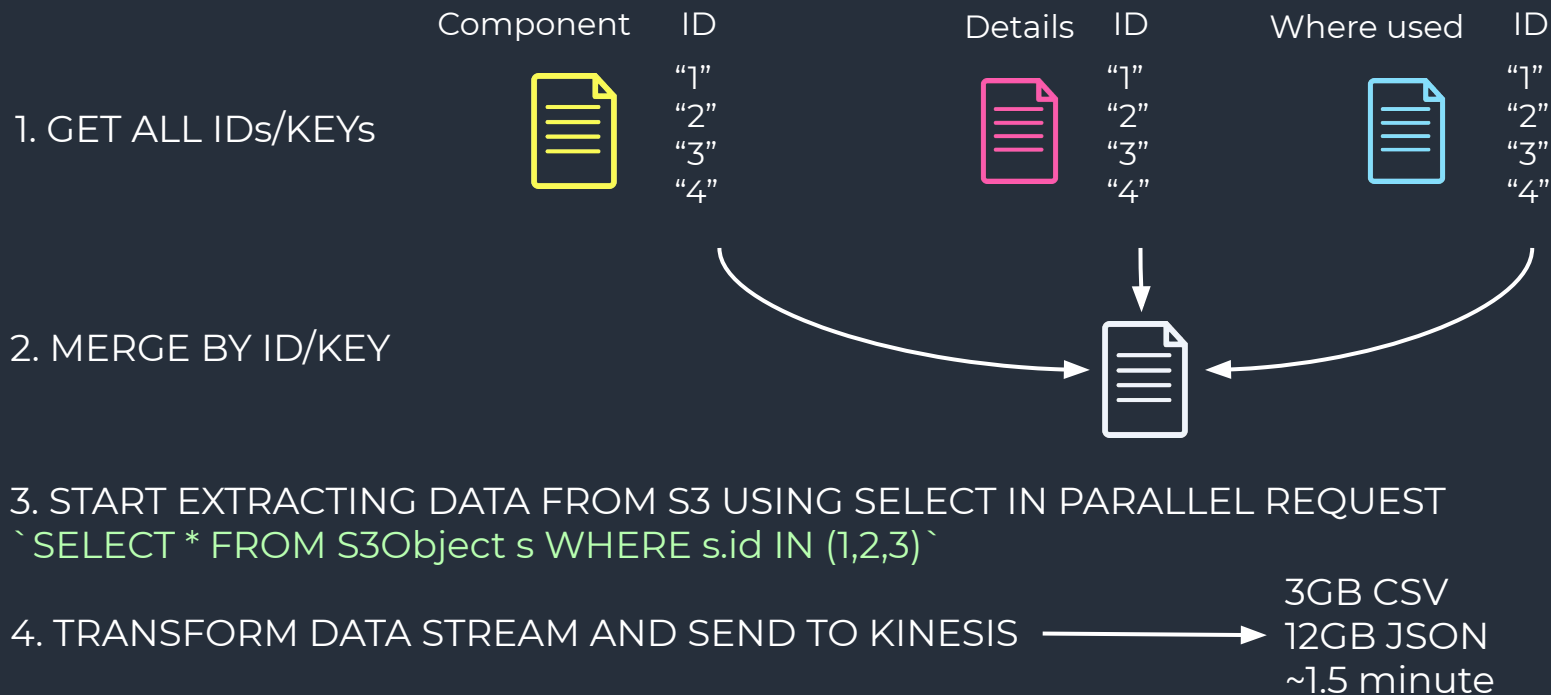
AGGREGATION

Component Details Where used

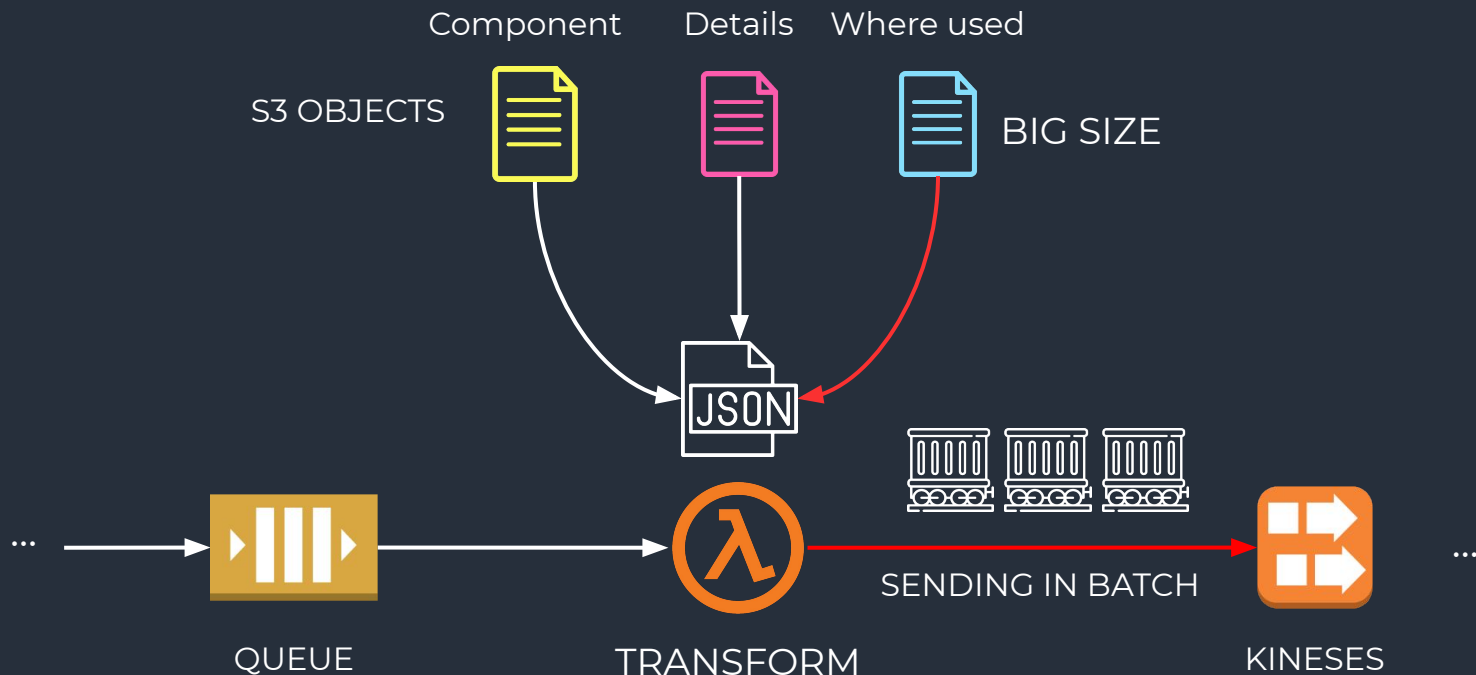


COMPOSITION

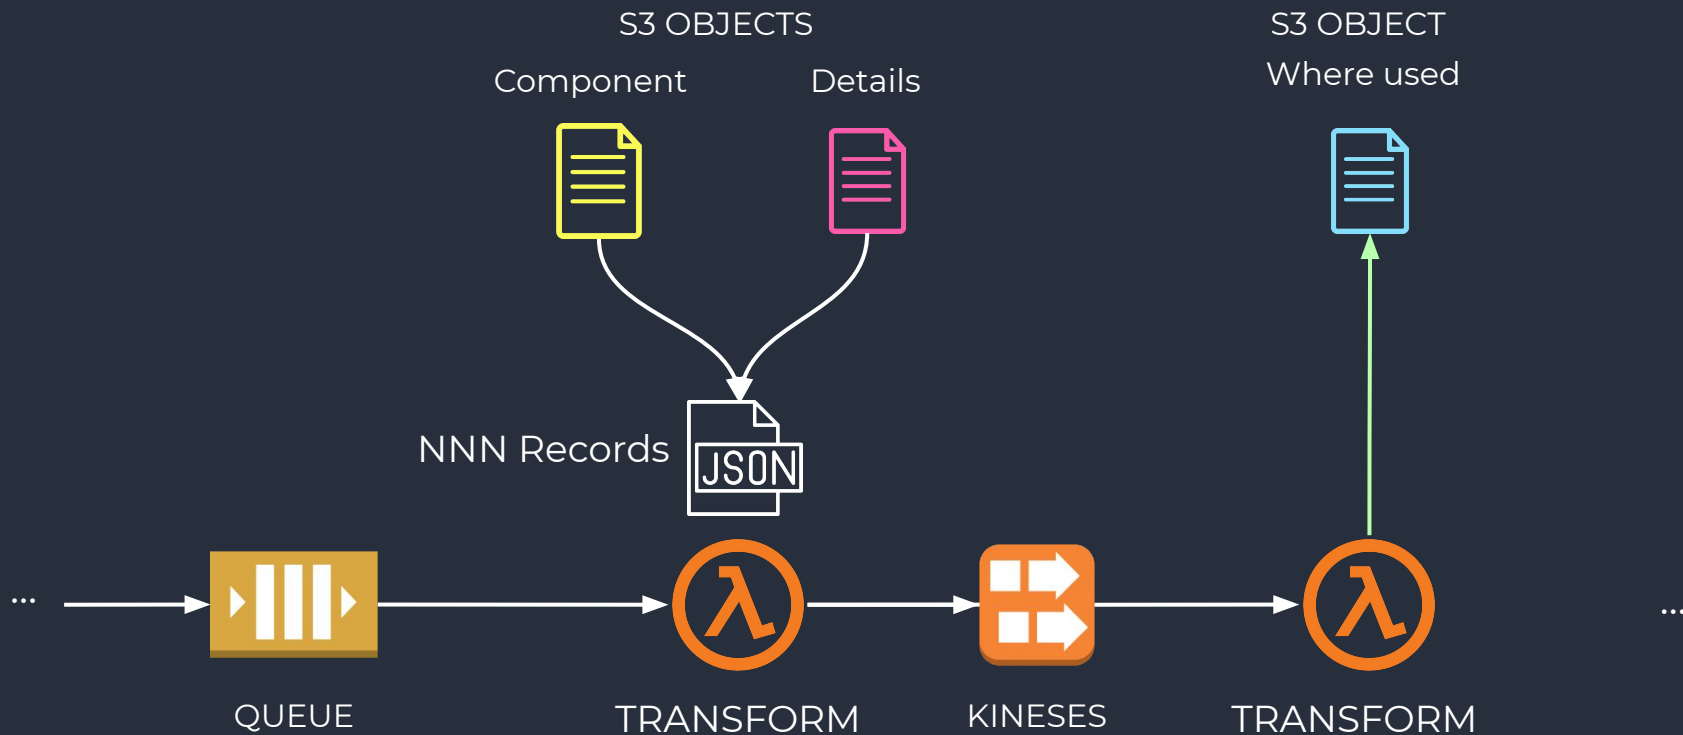
BIG FILES PROCESSING



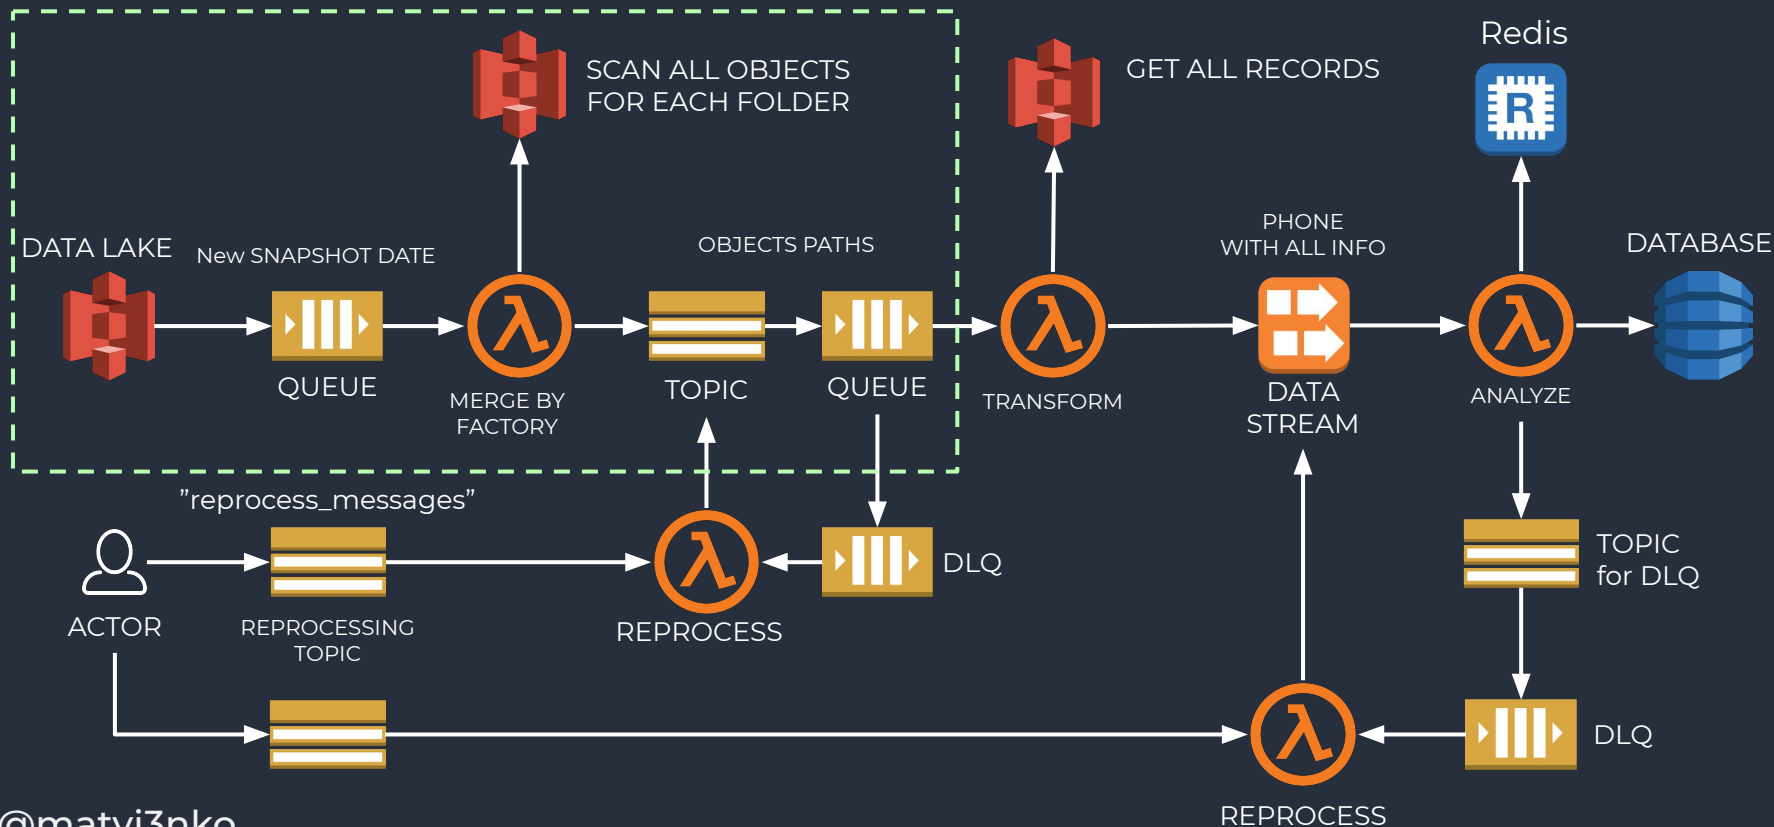
SENDING BIG OBJECTS



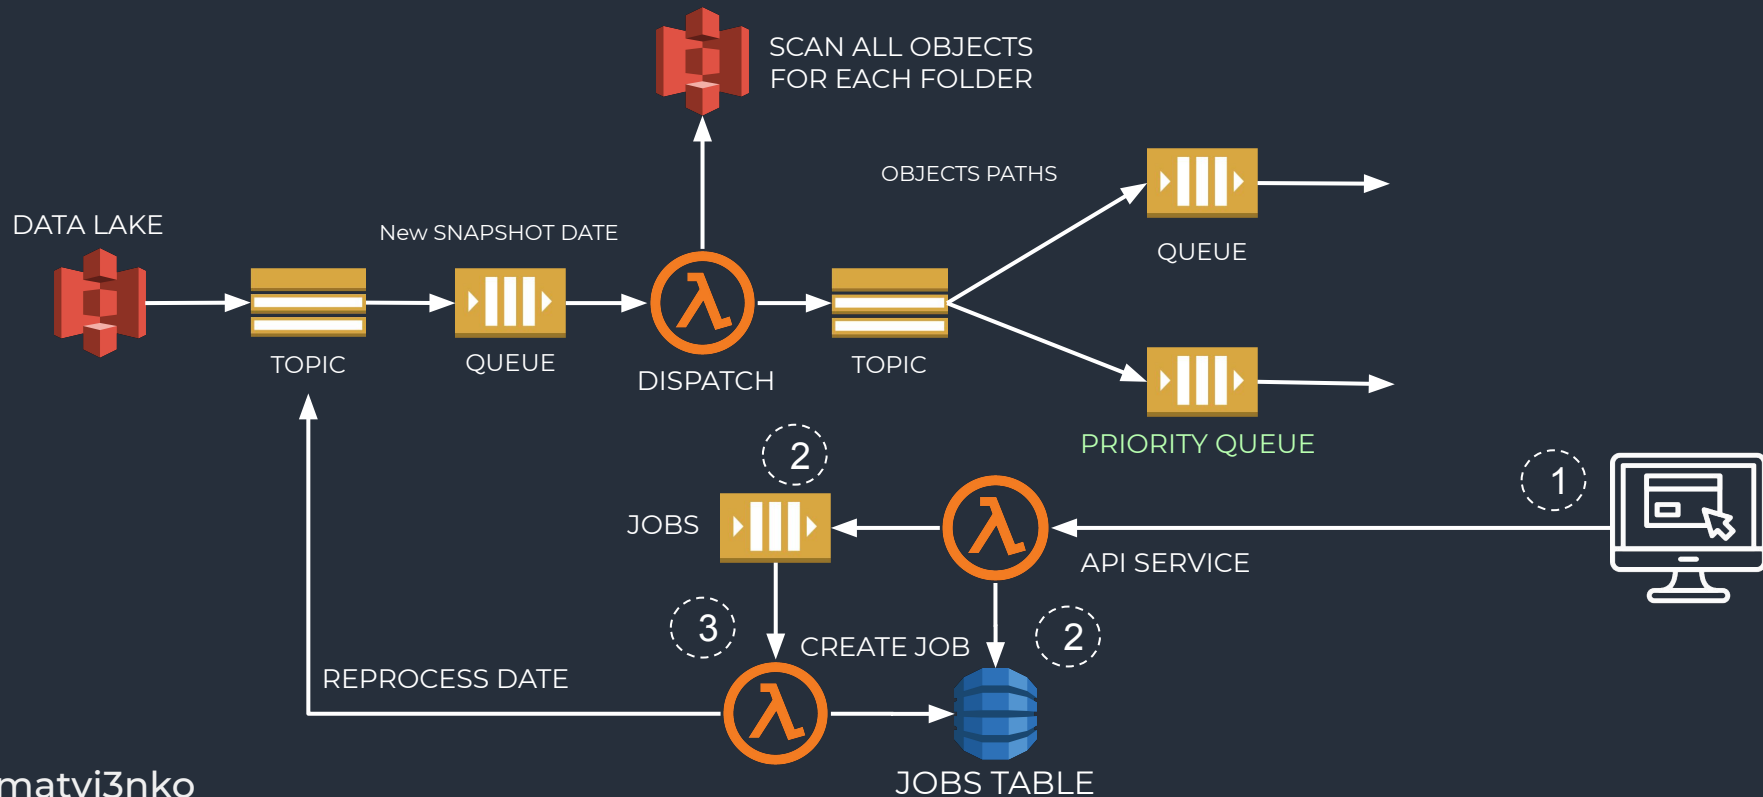
DECOUPLING



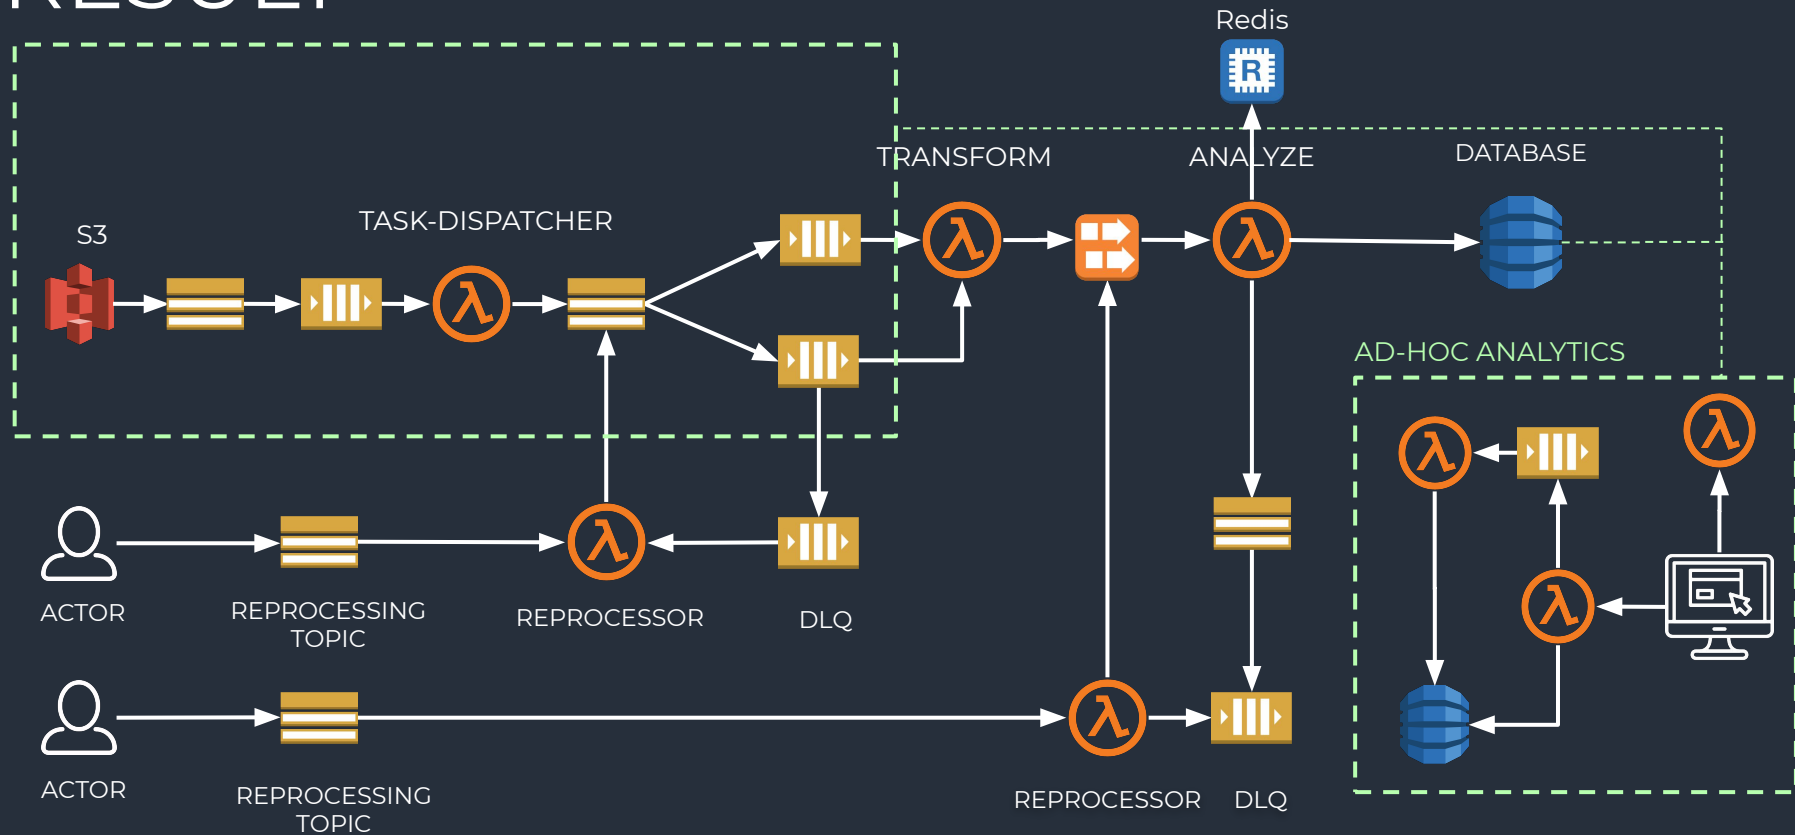
AHEAD OF THE QUEUE



ON-DEMAND REPROCESSING




RESULT



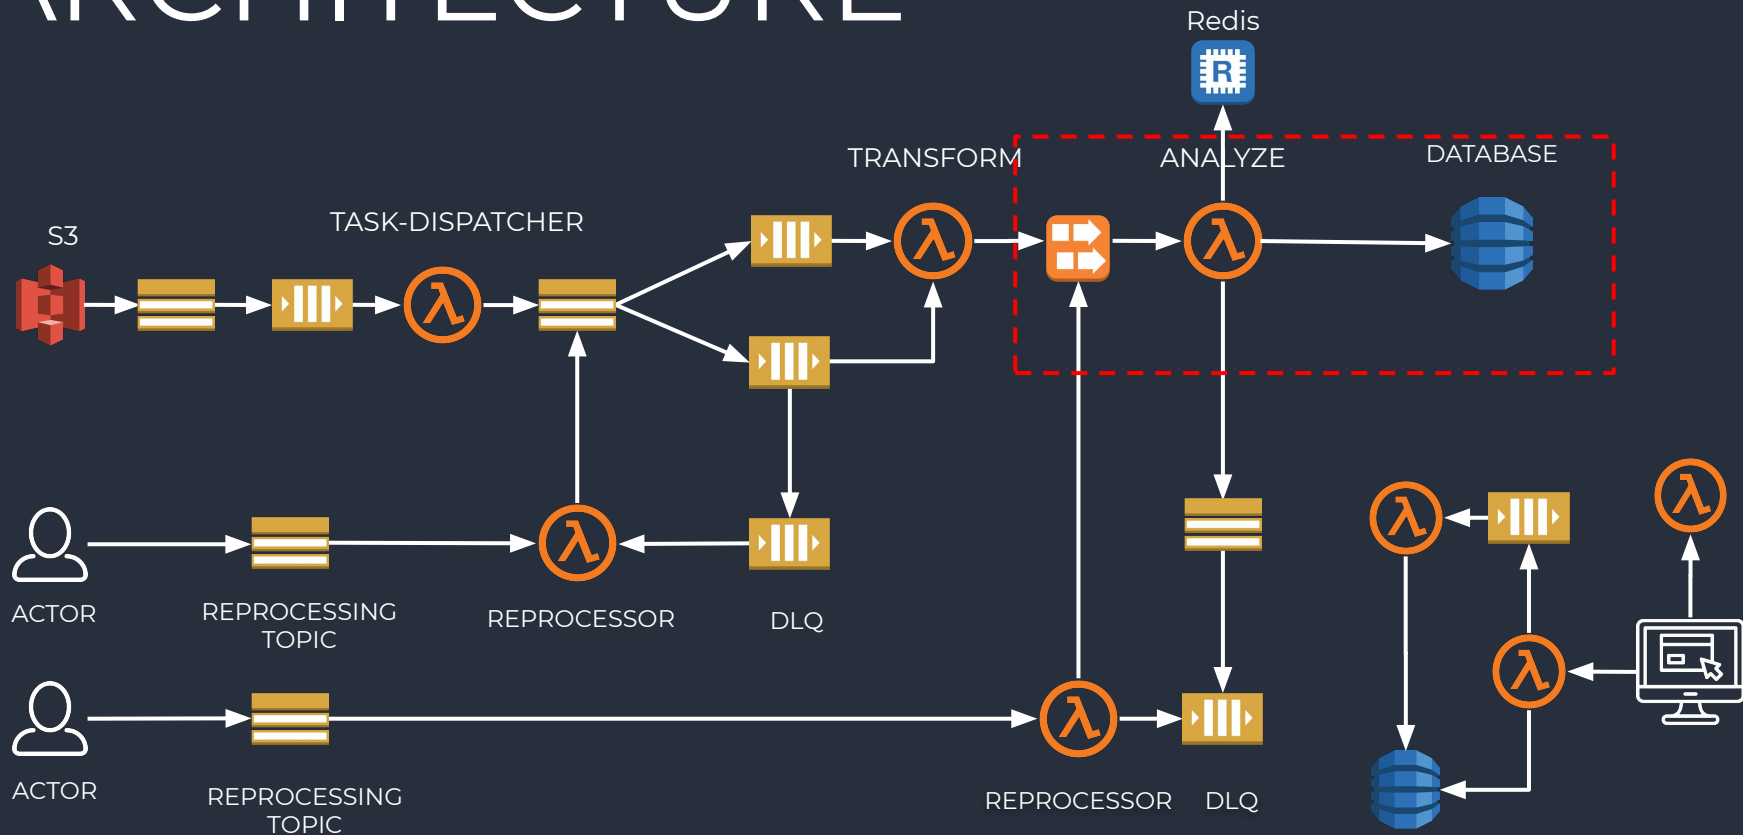
CONCLUSION

1. USE DATA COMPRESSION
2. GZIP RECORDS BEFORE PUT TO KINESIS STREAM
3. INCREASE MEMORY (CPU AUTO-LY) TO WORK WITH BIG OBJECTS/FILES
4. BUT DON'T BUFFER RESPONSE, WORK WITH STREAMS


PROGRESS




ARCHITECTURE



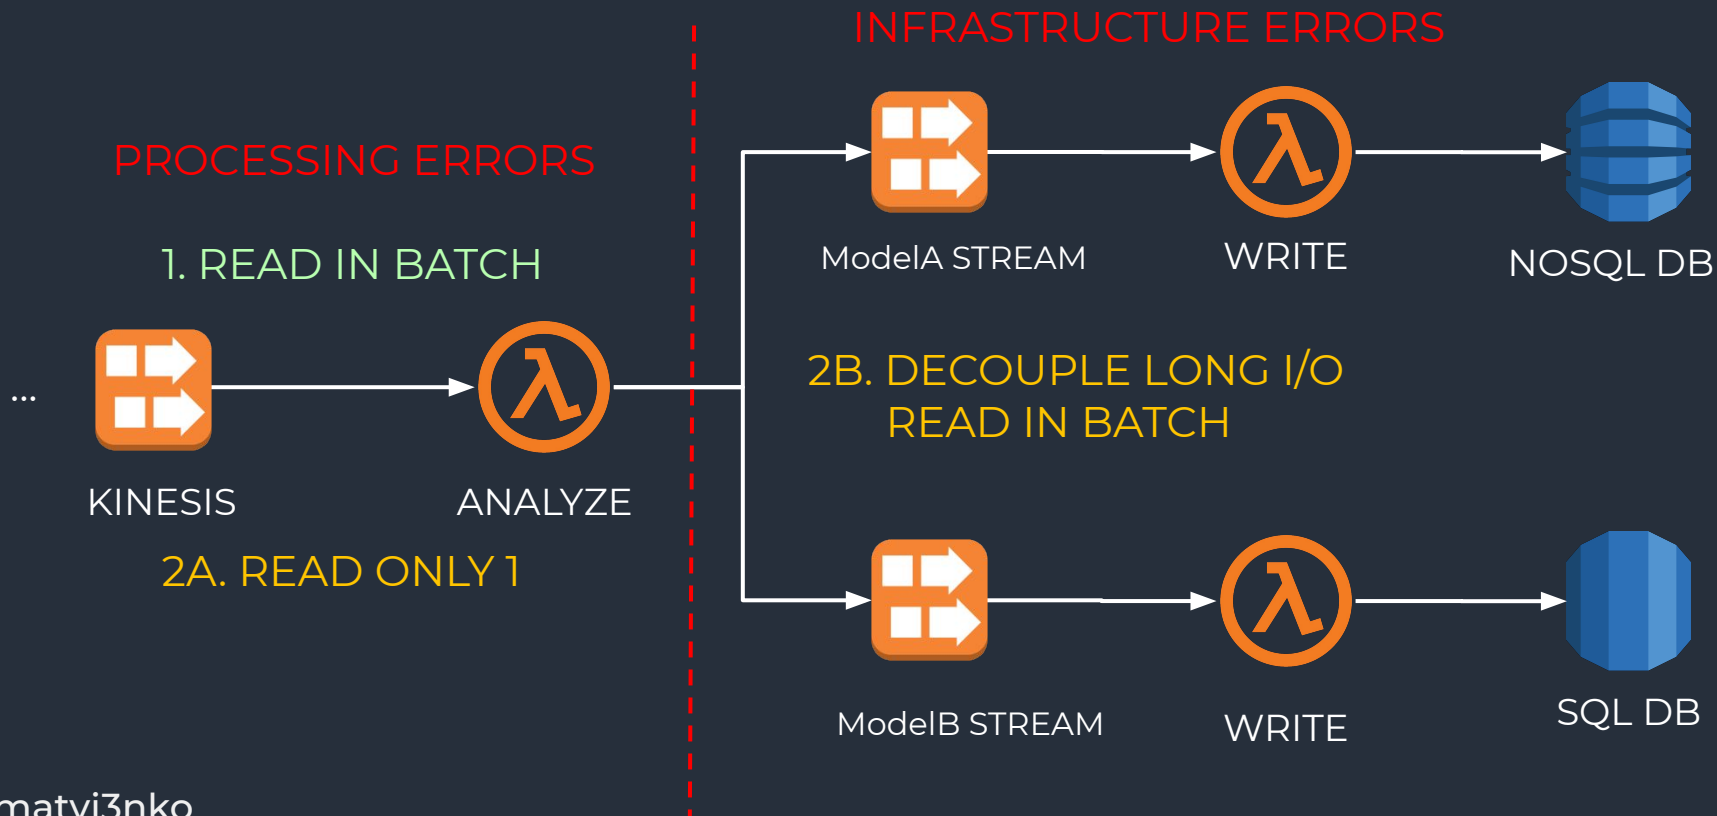
PROBLEM



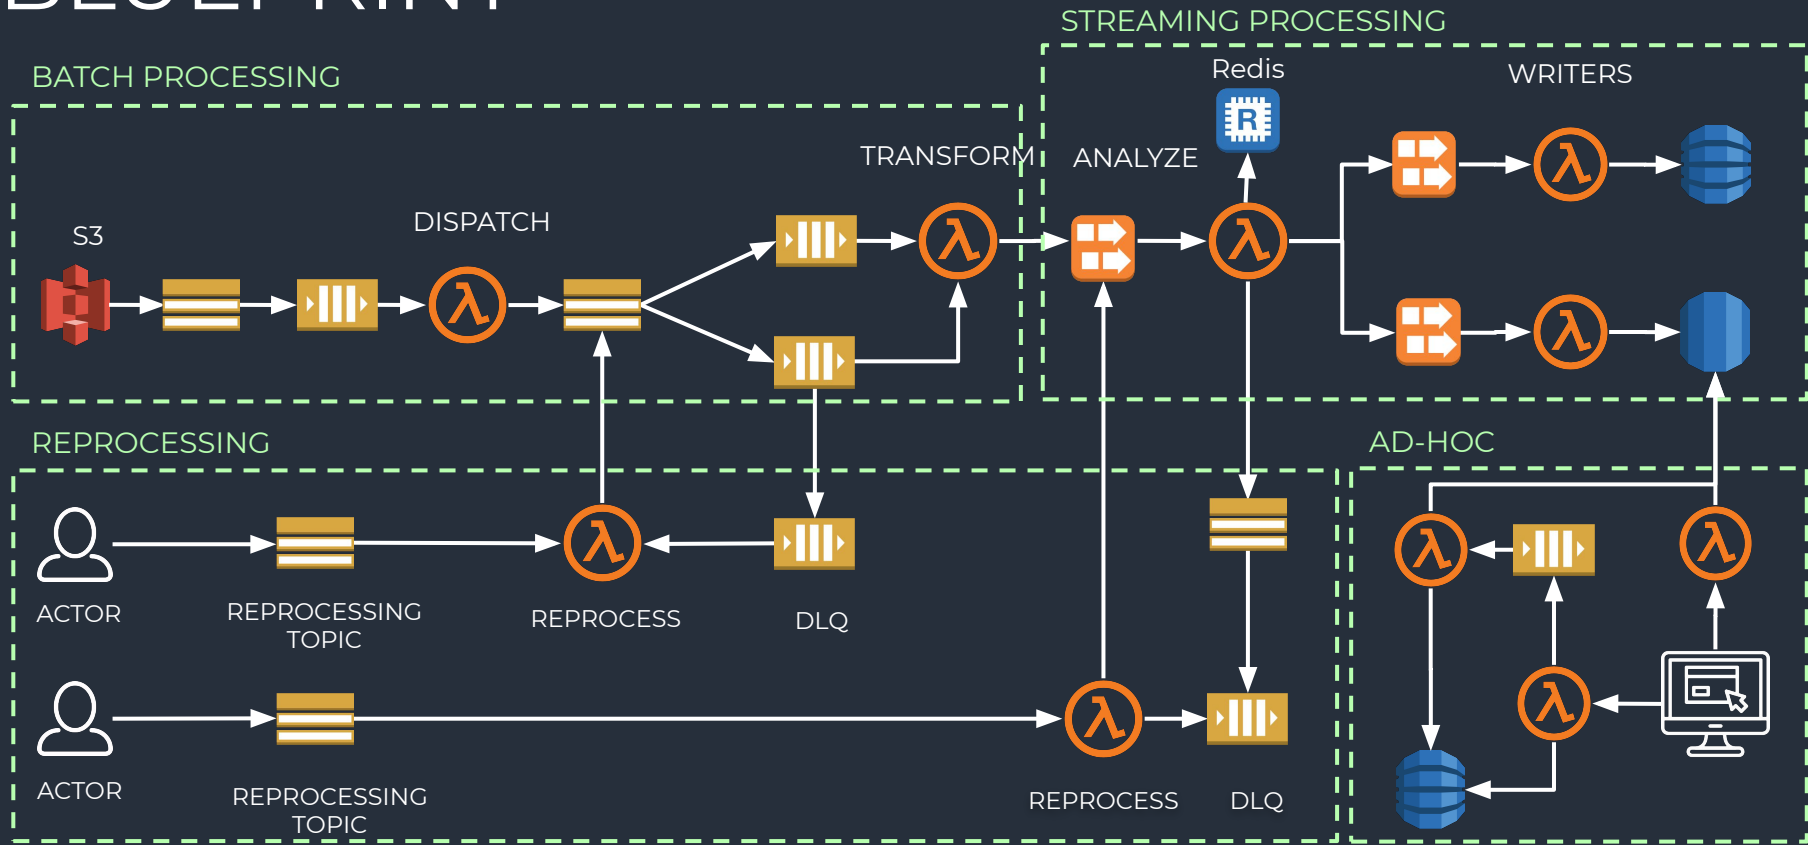
SHARDS OF THE STREAM




DECOUPLING



BLUEPRINT



PROGRESS



SCALABILITY

READ AND PROCESS DATA

ERROR HANDLING

SOLUTION STRUCTURE AND FUNCTION

BASE ARCHITECTURE DESIGN

CASE STUDY

1. 25% to develop – 75% spend on integration, but more flexible for changes in a result.
2. 5K, 7K, **10K lambda functions in parallel**
3. Terabytes of data
4. Serverless vs EC2 cost

SUMMARY

1. PERFORMANCE
2. TROUBLESHOOTING
3. MEMORY LEAKS
4. PROVIDER LOCK
5. ARCHITECTURE

THANKS!



Nikolay Matvienko

matvi3nko@gmail.com

Twitter.com/matvi3nko

github.com/matvi3nko