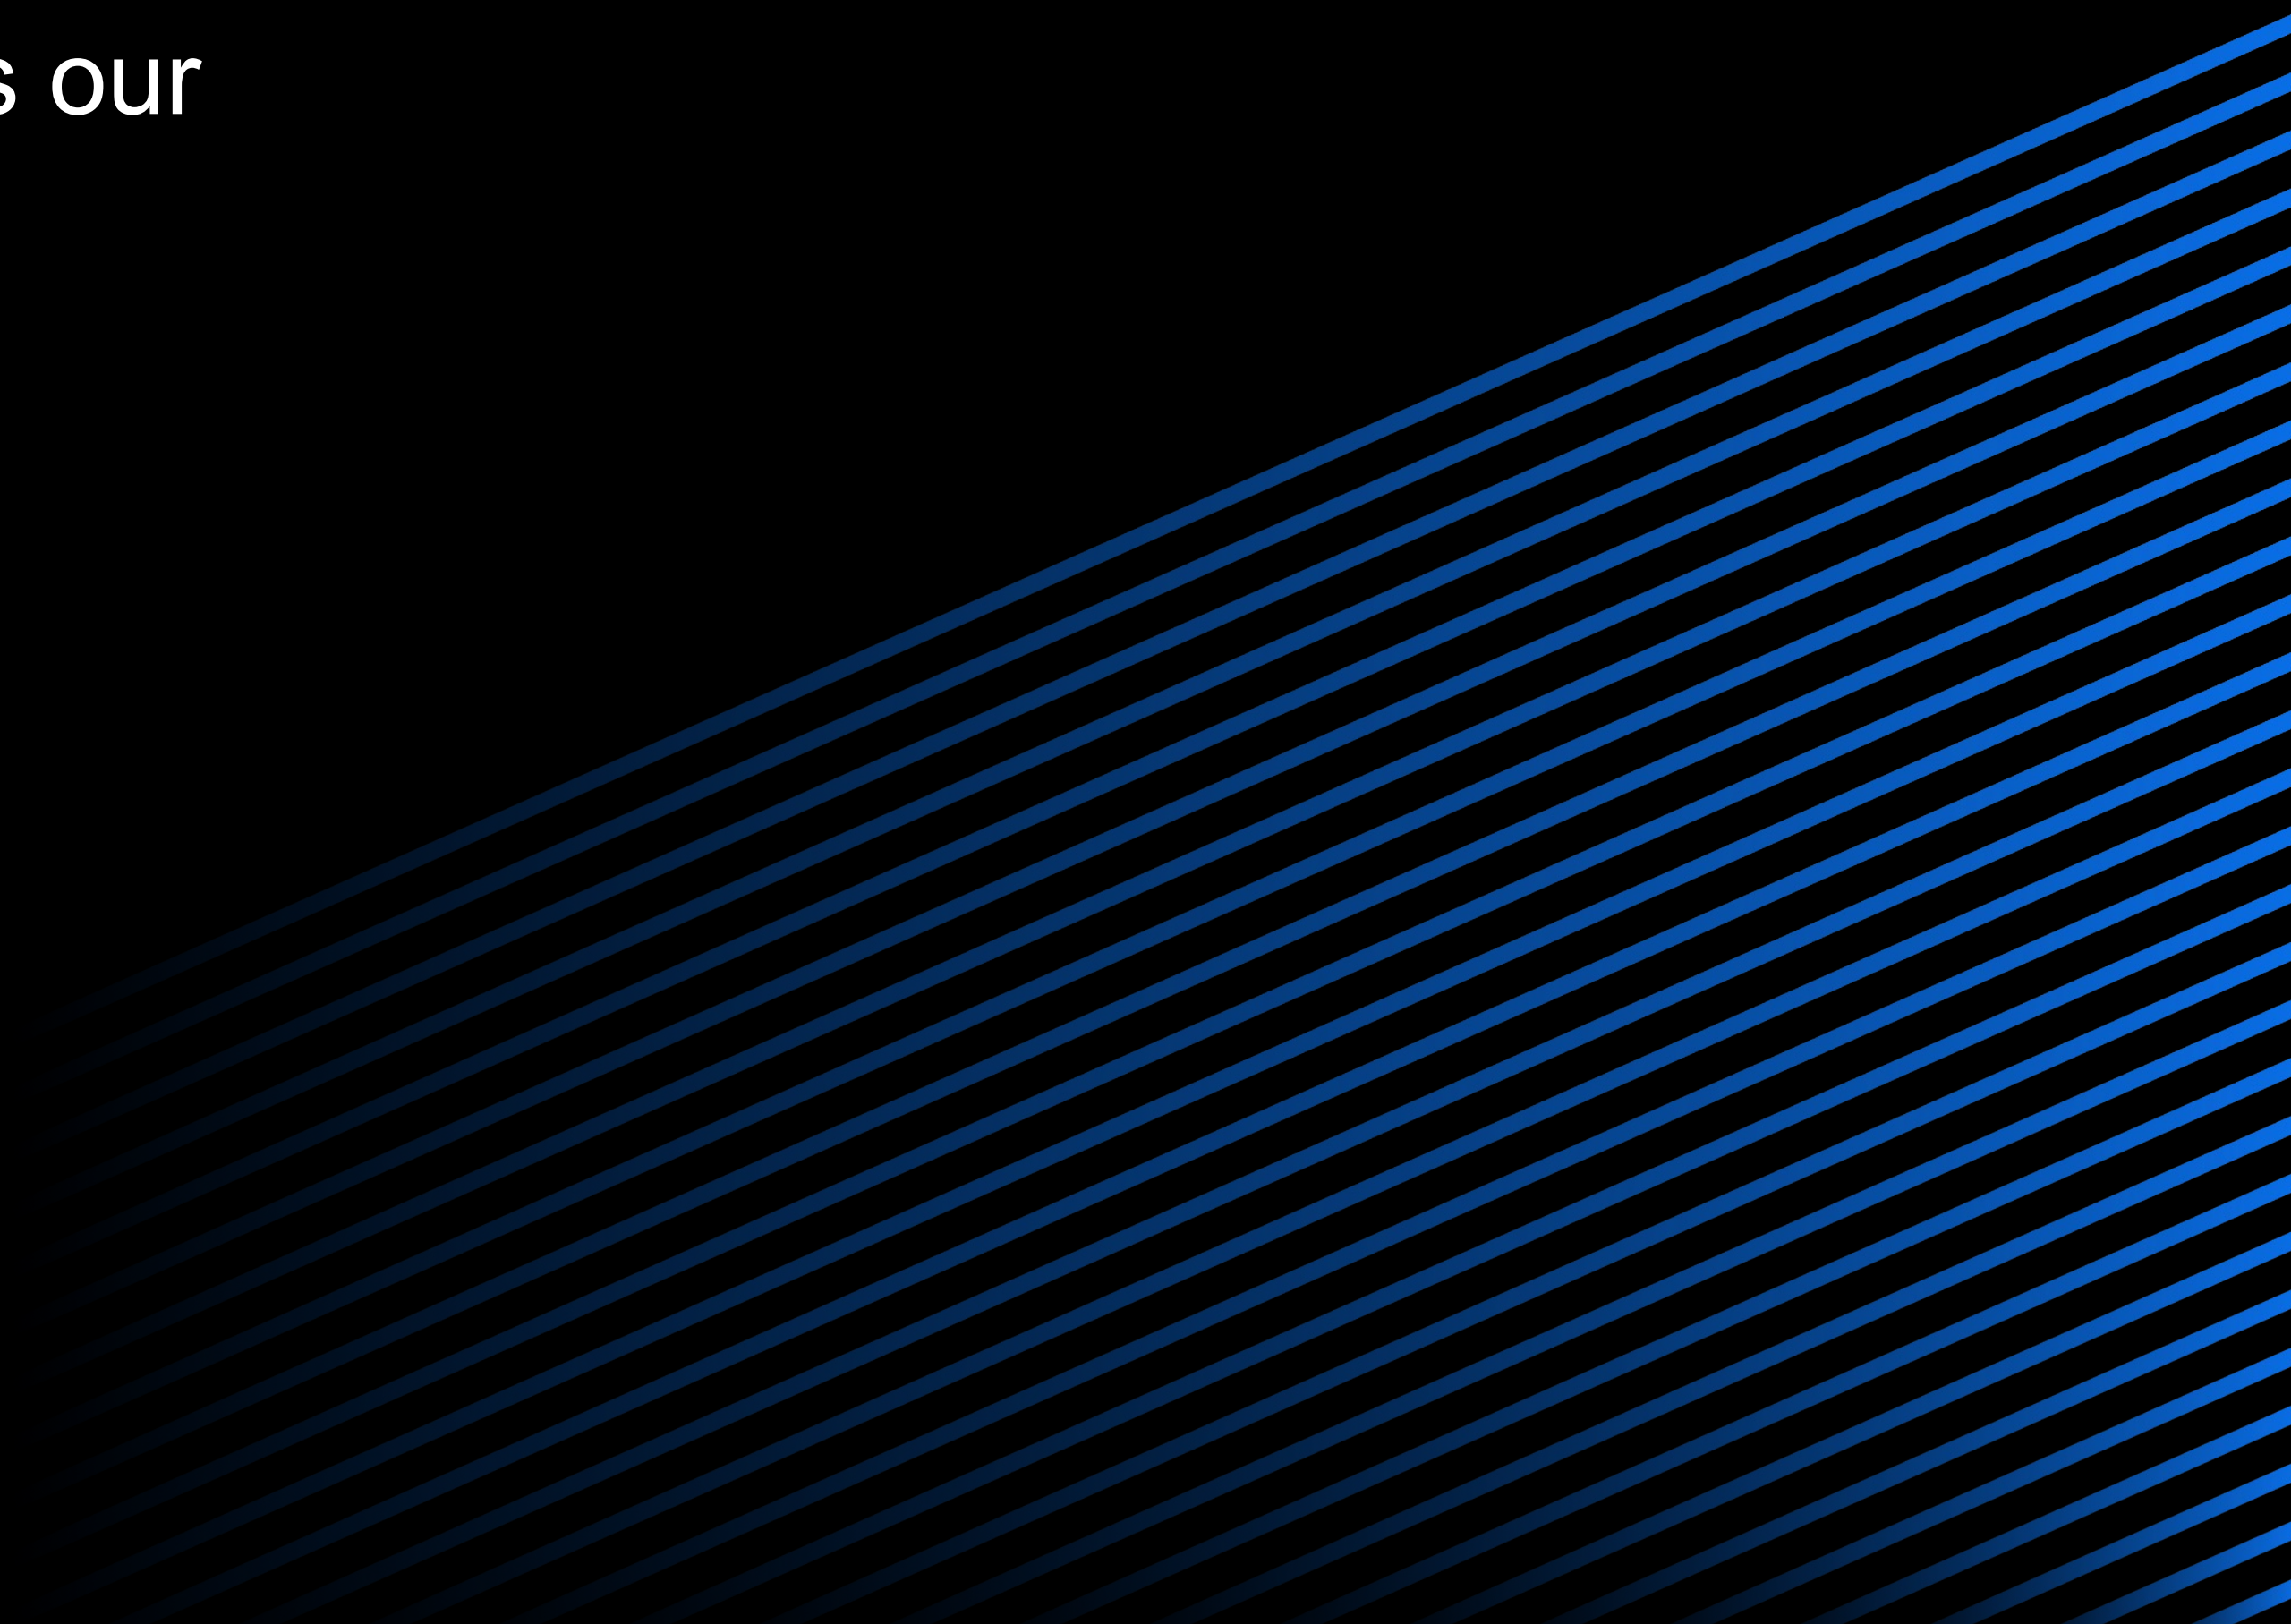


# LINE TODAY - 微服務架構支撐千萬 活躍用戶的影音內容平臺

Libra Huang - LINE TODAY

10/18/2022

# Agenda

- LINE TODAY and its architecture
  - How miniservices and K8S changes our development and operation
  - Lessons learnt
- 

# LINE TODAY Stays With You Today



個人化訊息

8:00

財經天氣

生活內容

15:00

話題/投票/電影

直播

20:00


棒球/NBA/演唱會

12:30

國際/國內/娛樂  
新聞內容

18:00

TODAY看世界  
官方帳號



**Taiwan**



**Thailand**



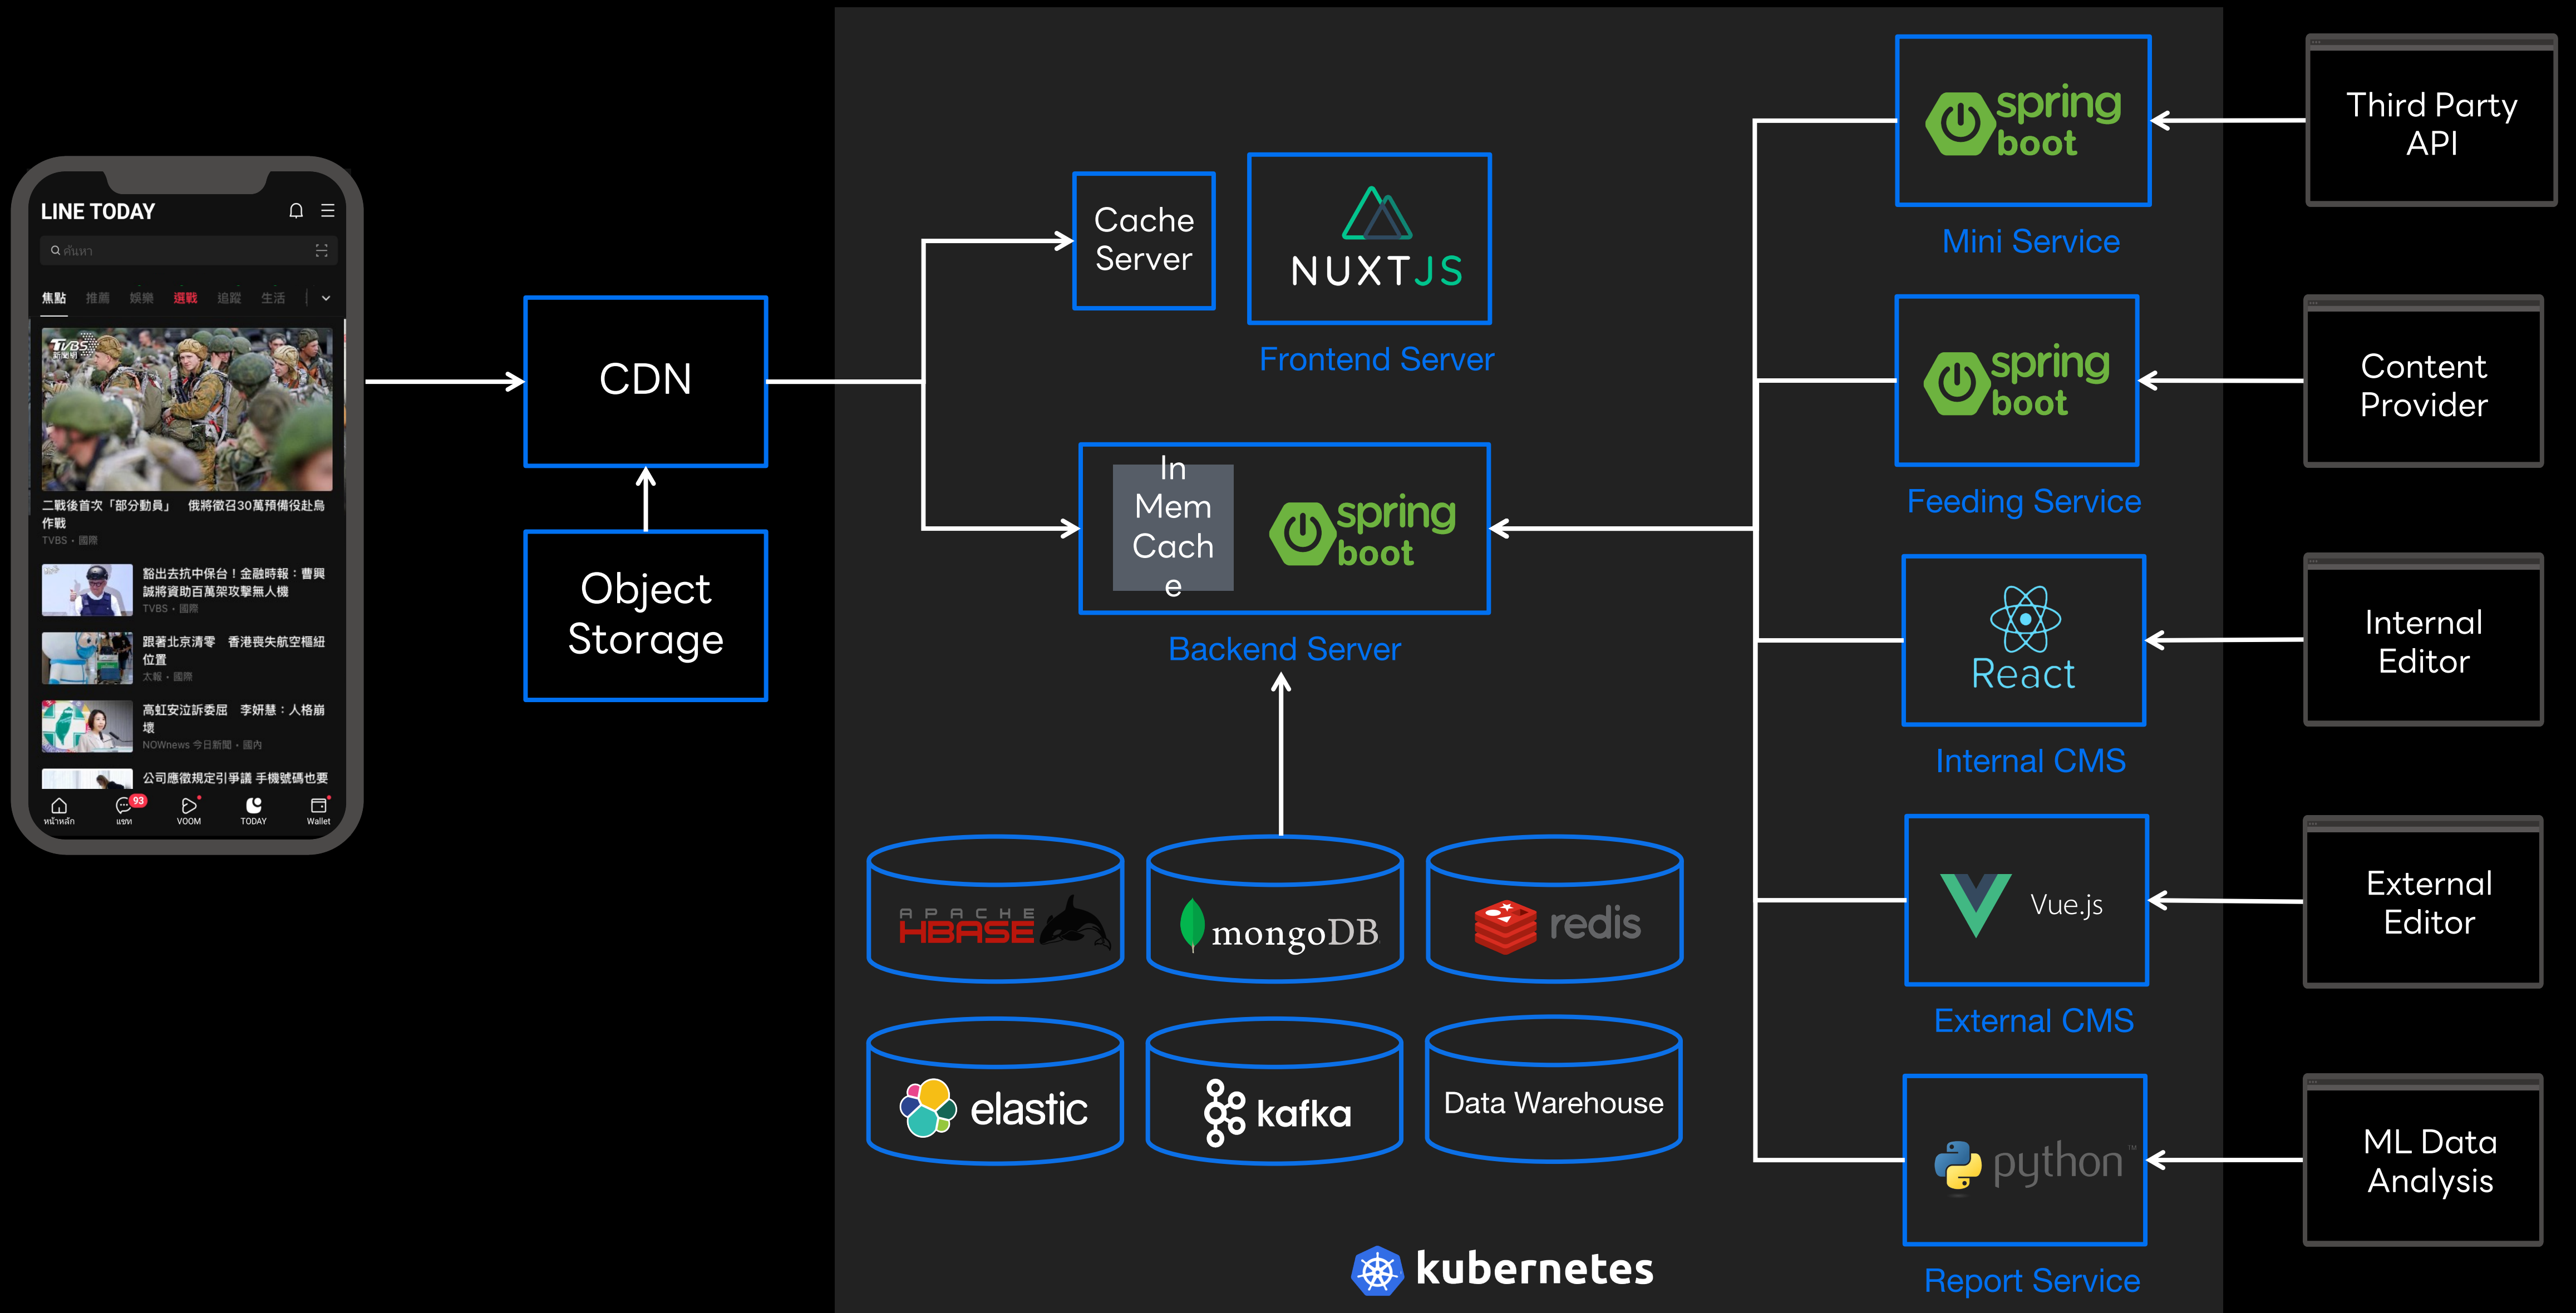
**Hong Kong**



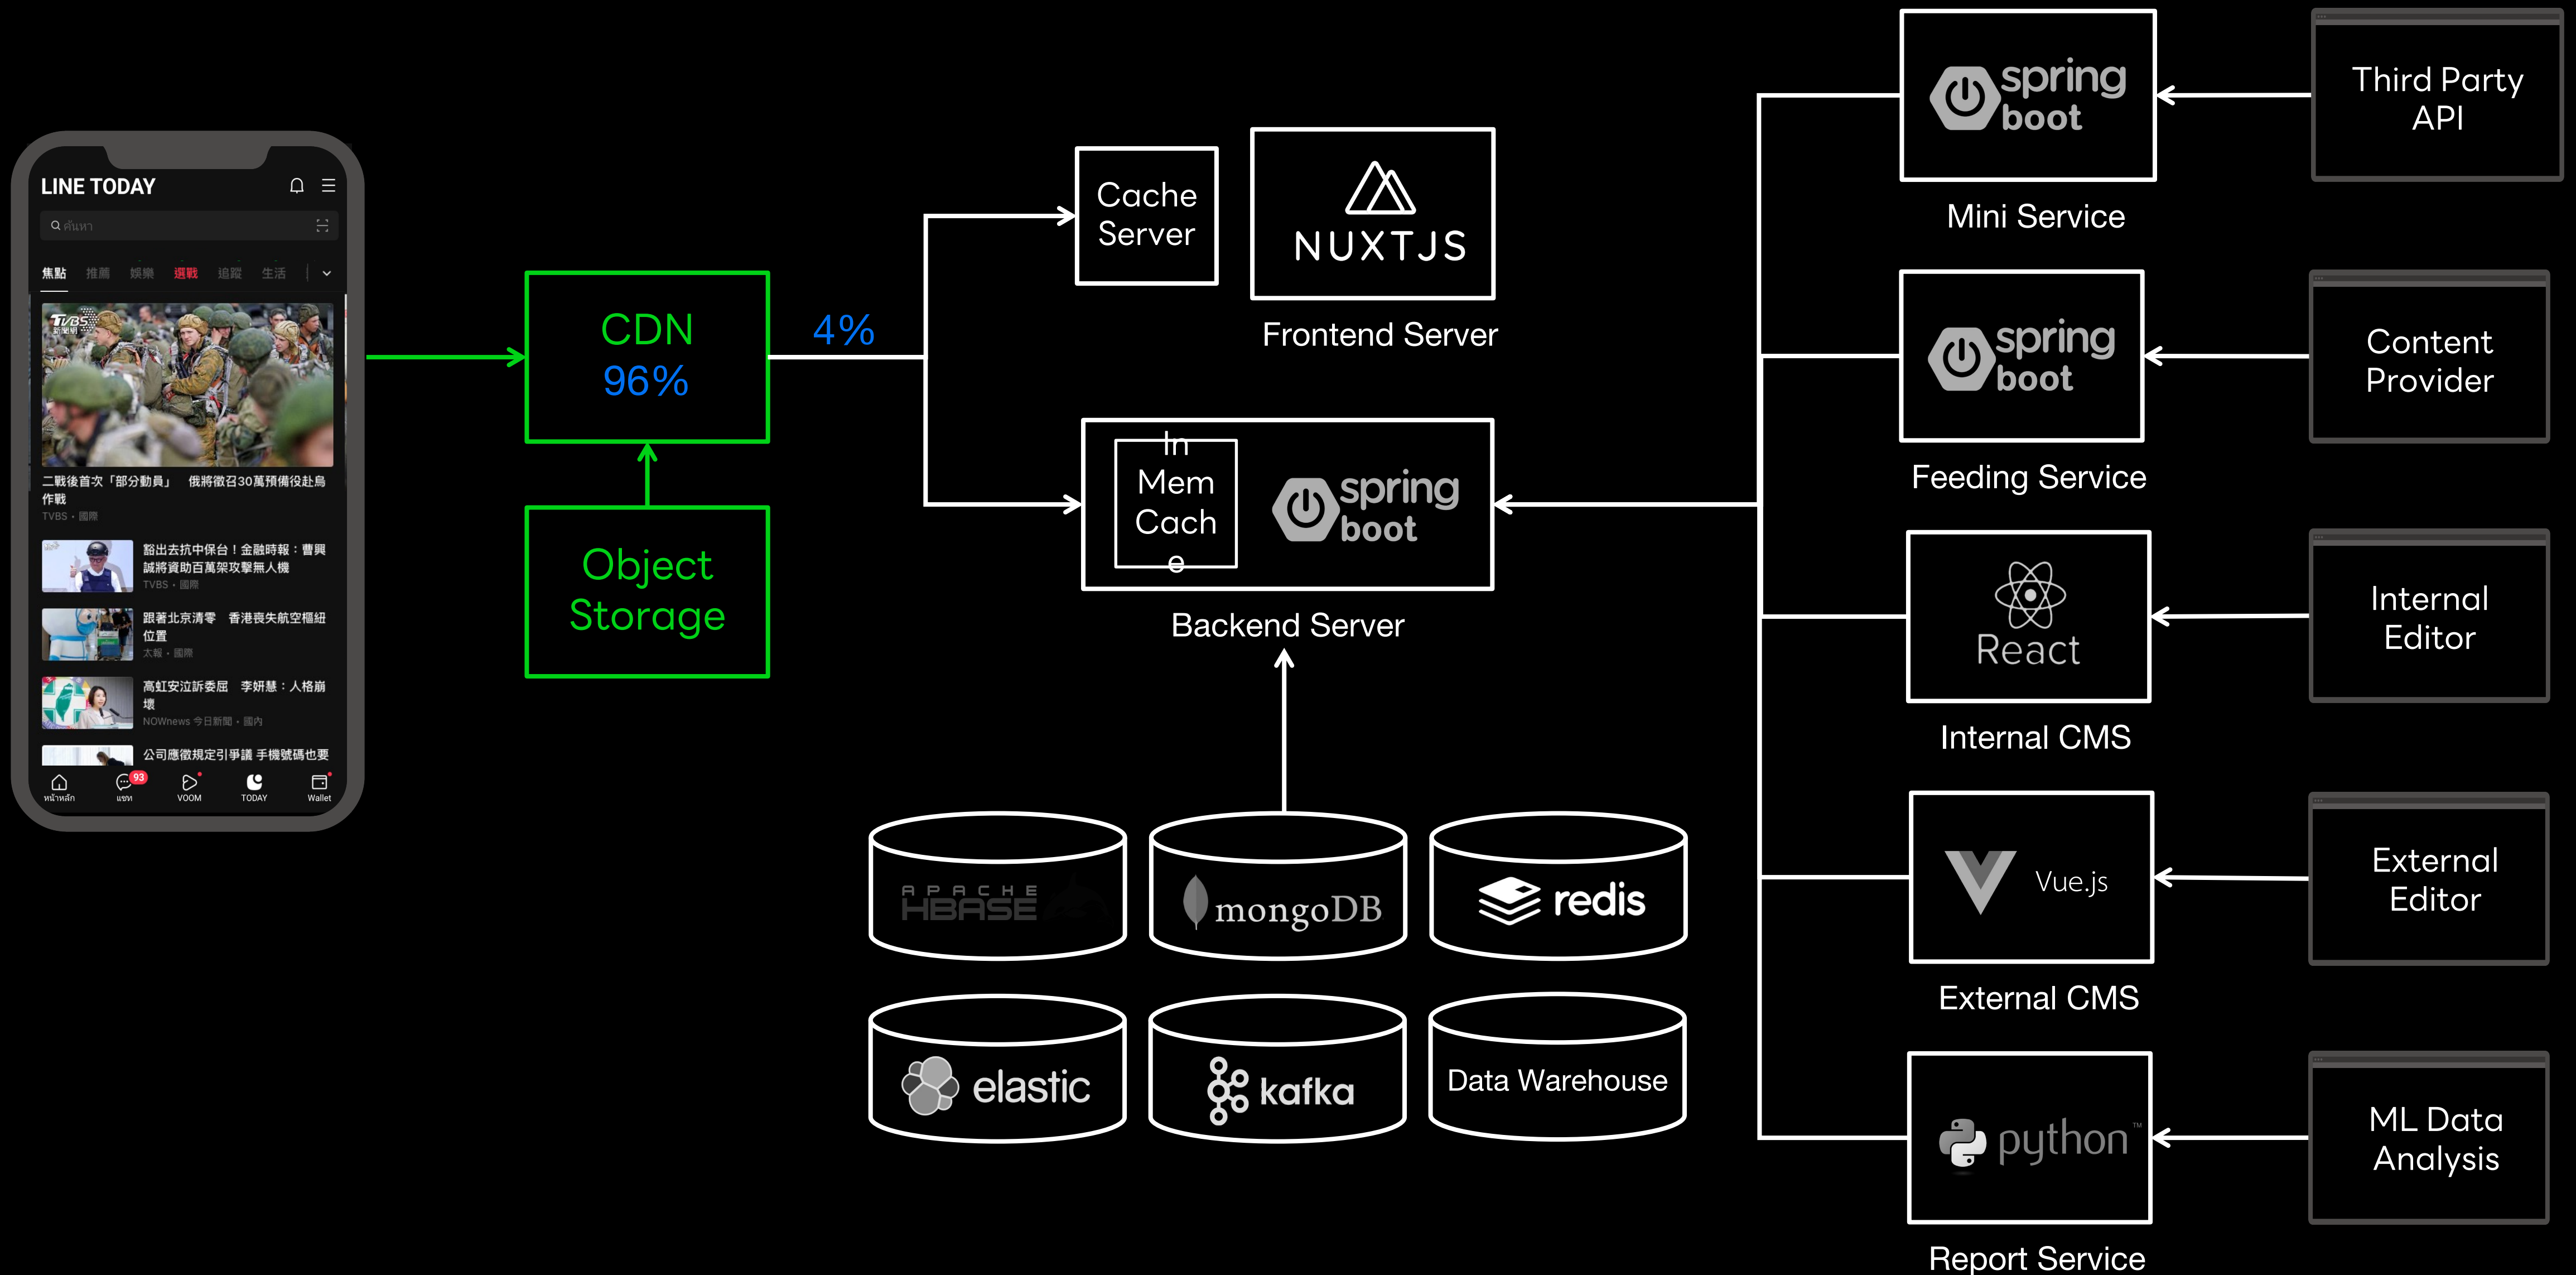
**18M MAU**

LINE TODAY Taiwan

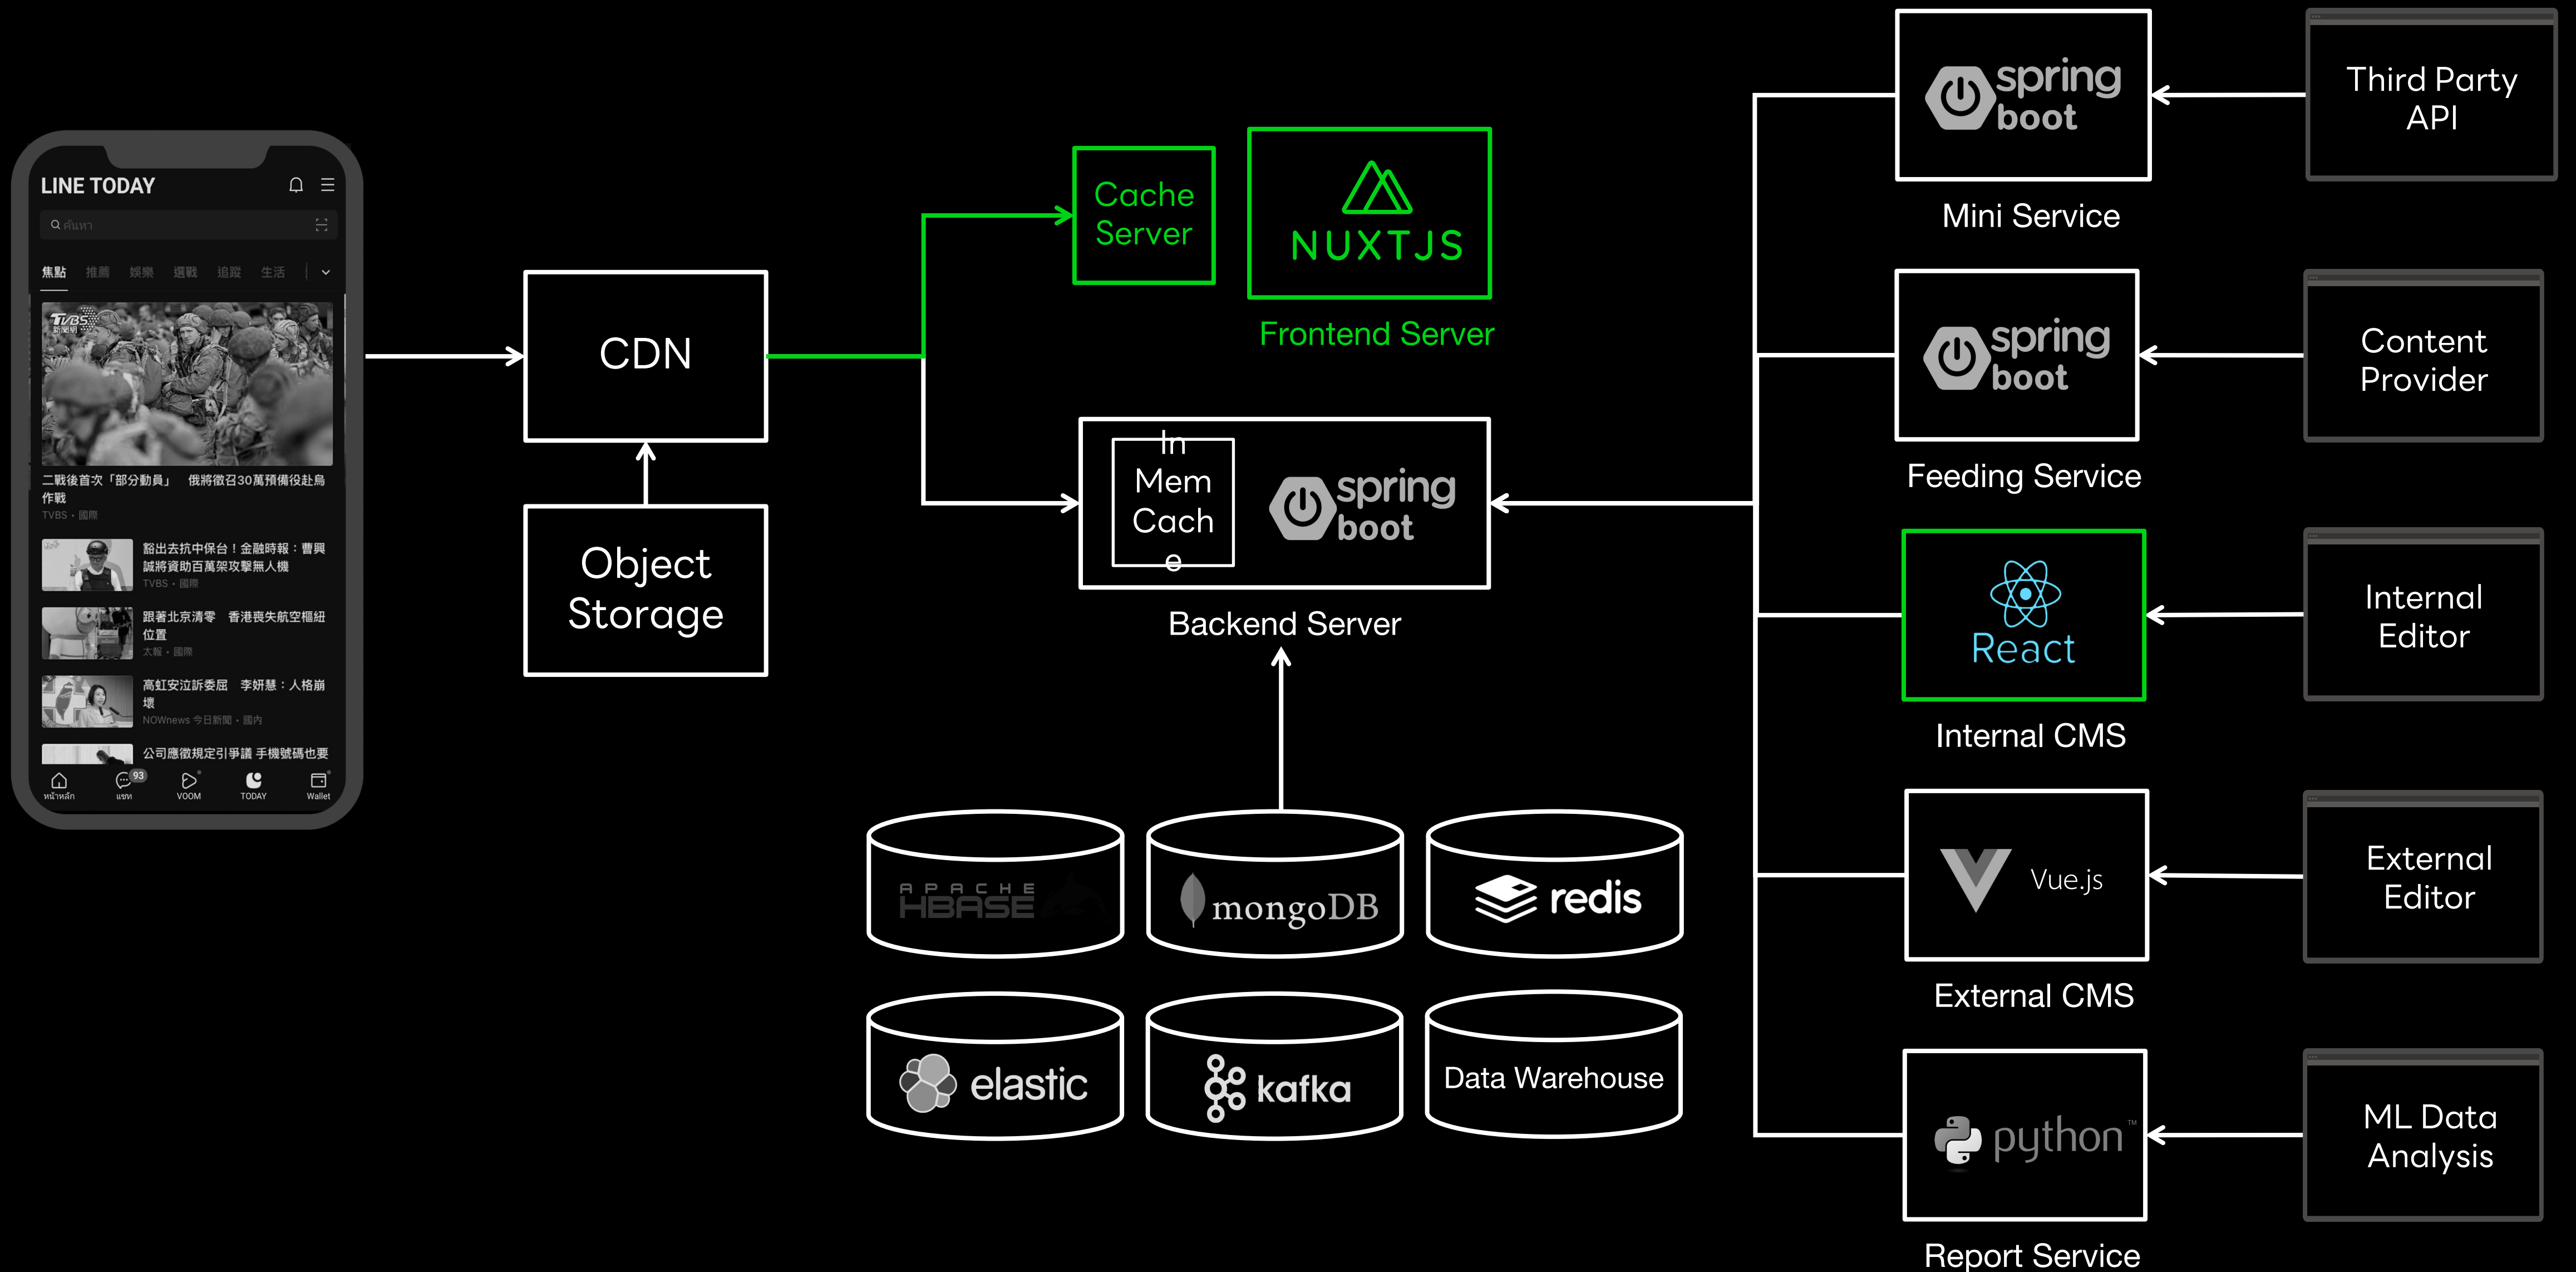
# LINE TODAY Architecture



# LINE TODAY Architecture

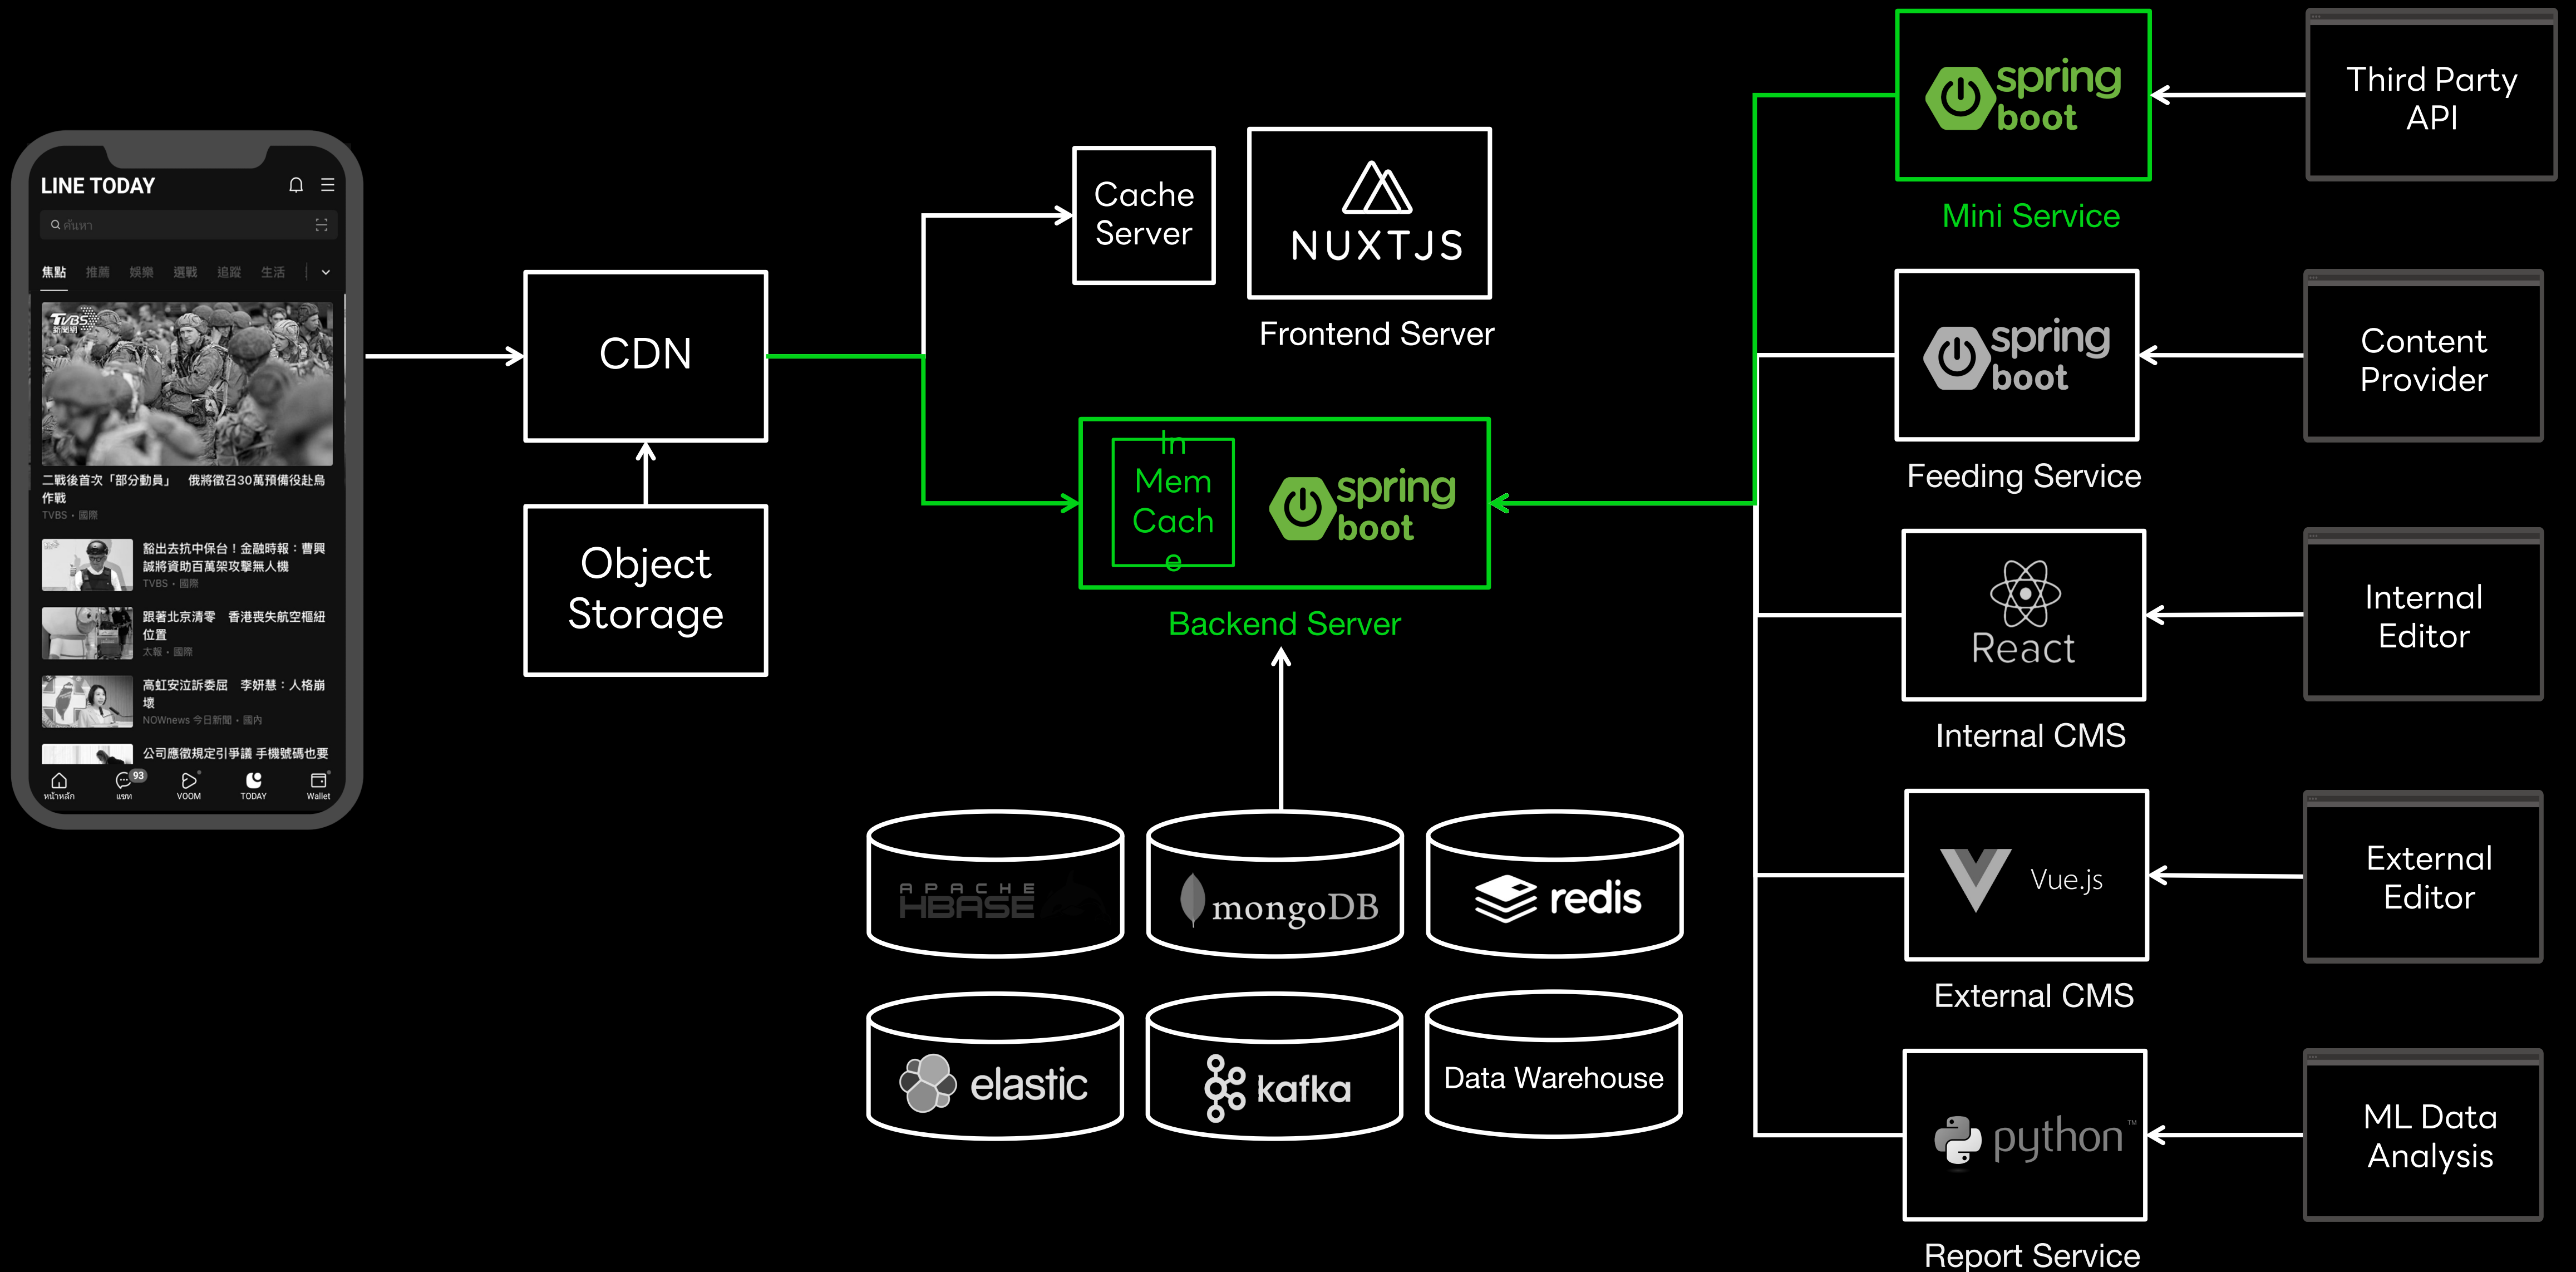


# LINE TODAY Architecture

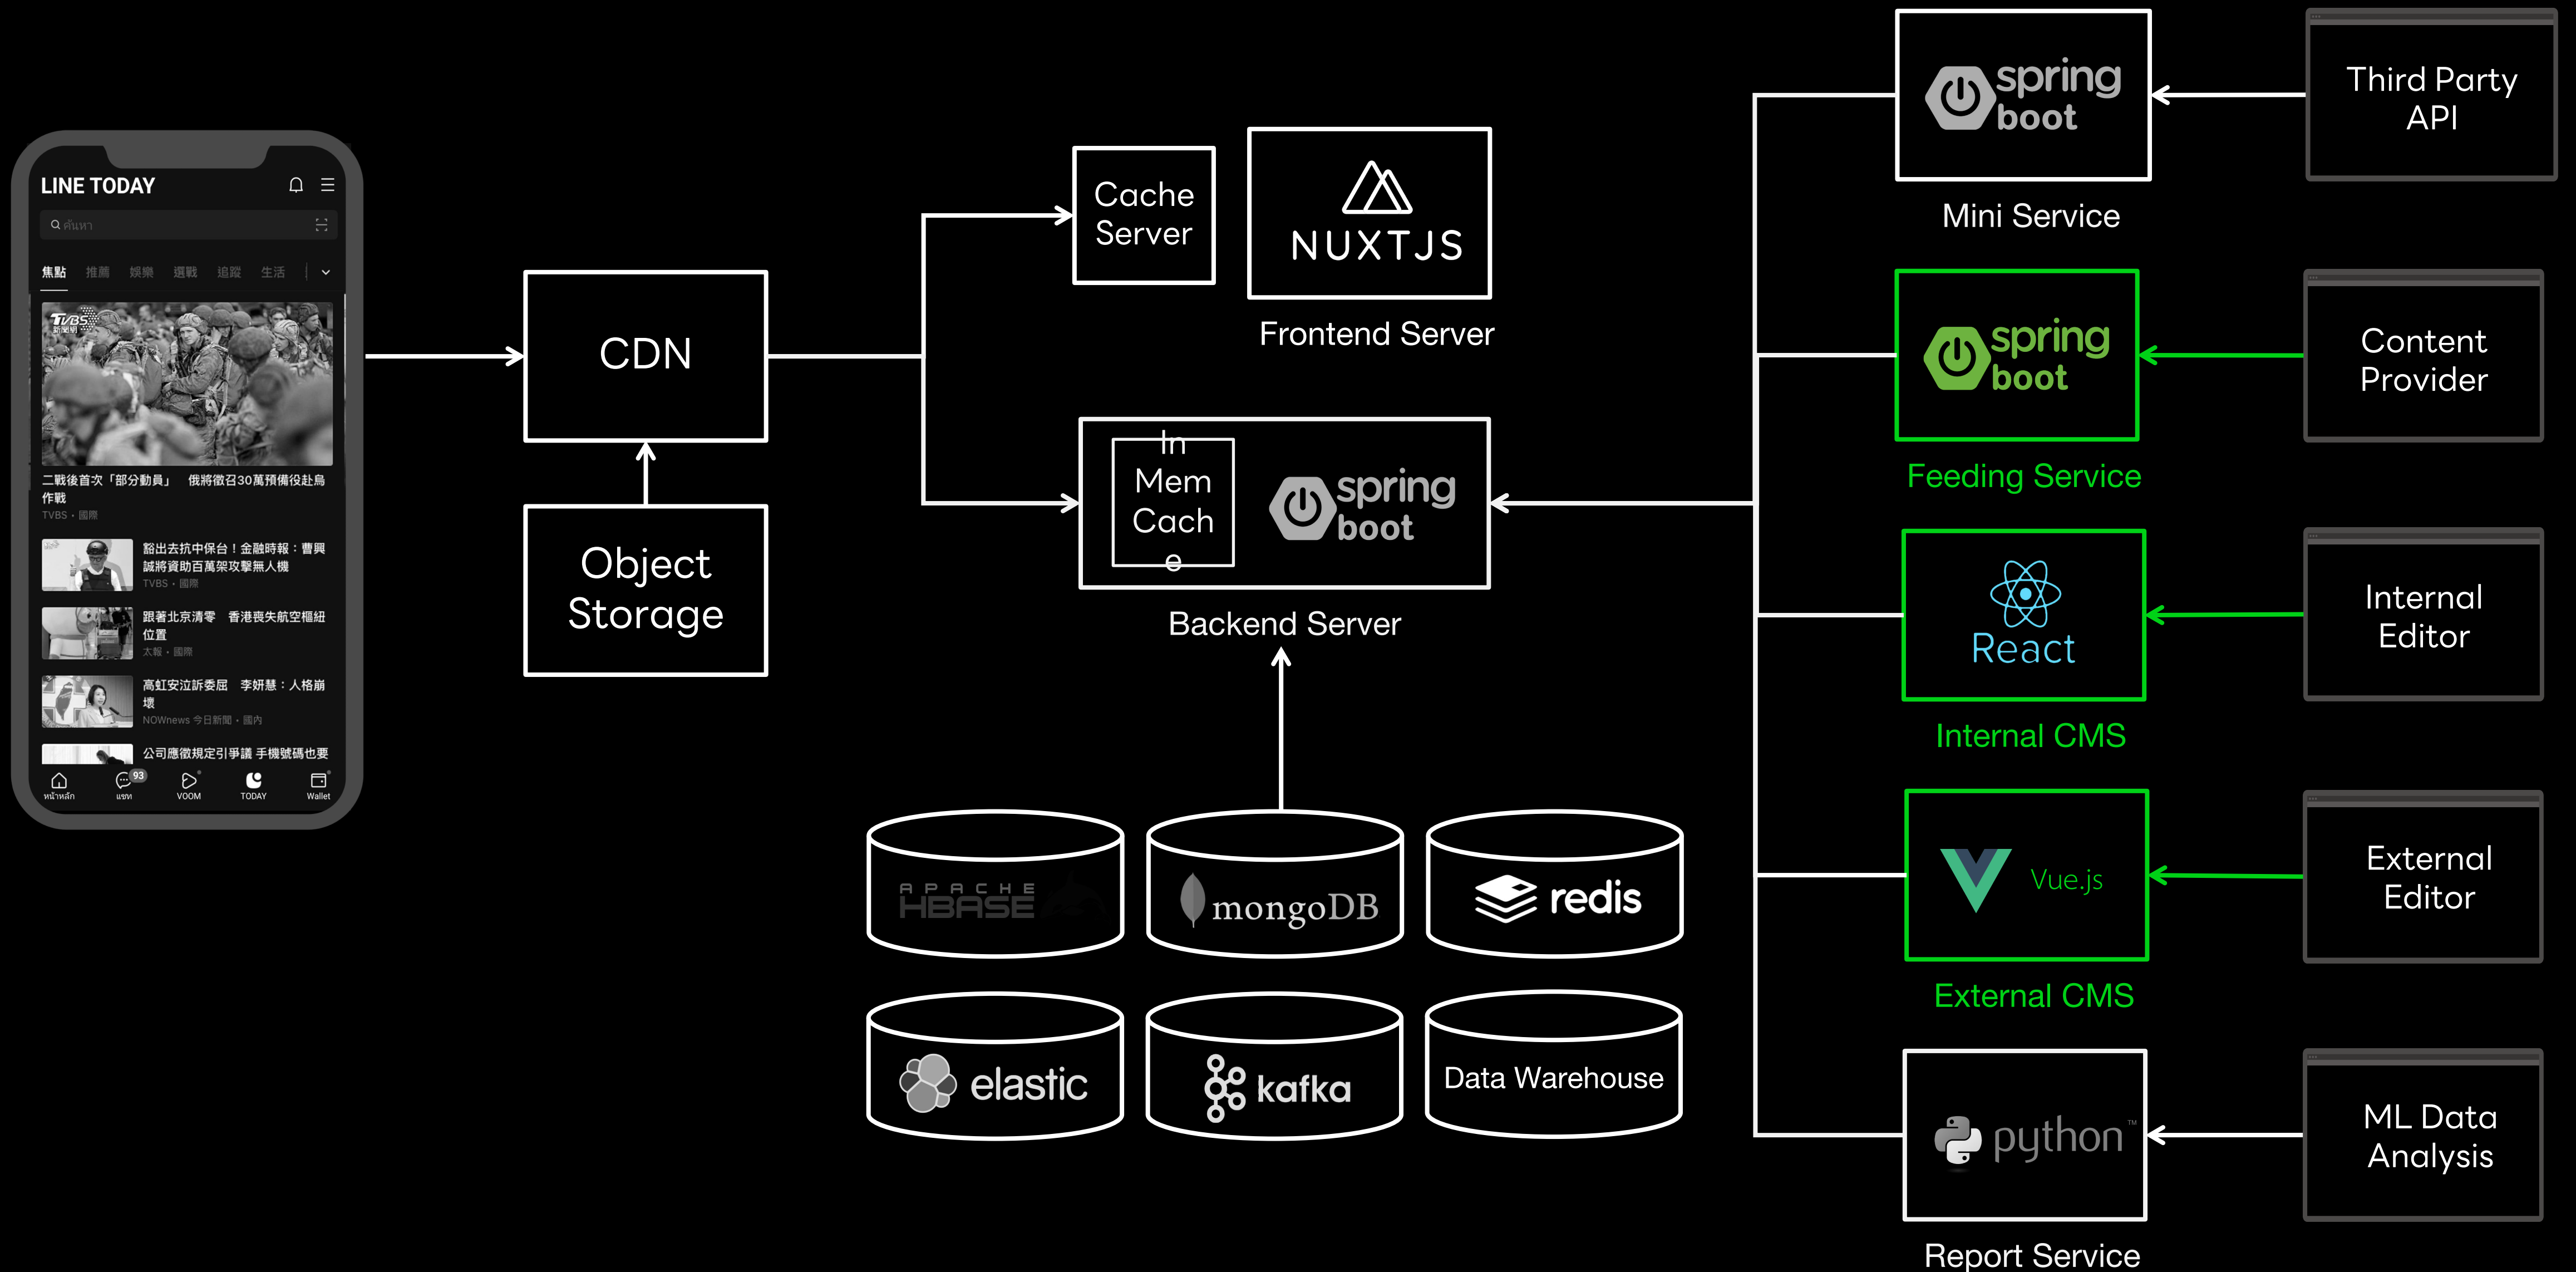




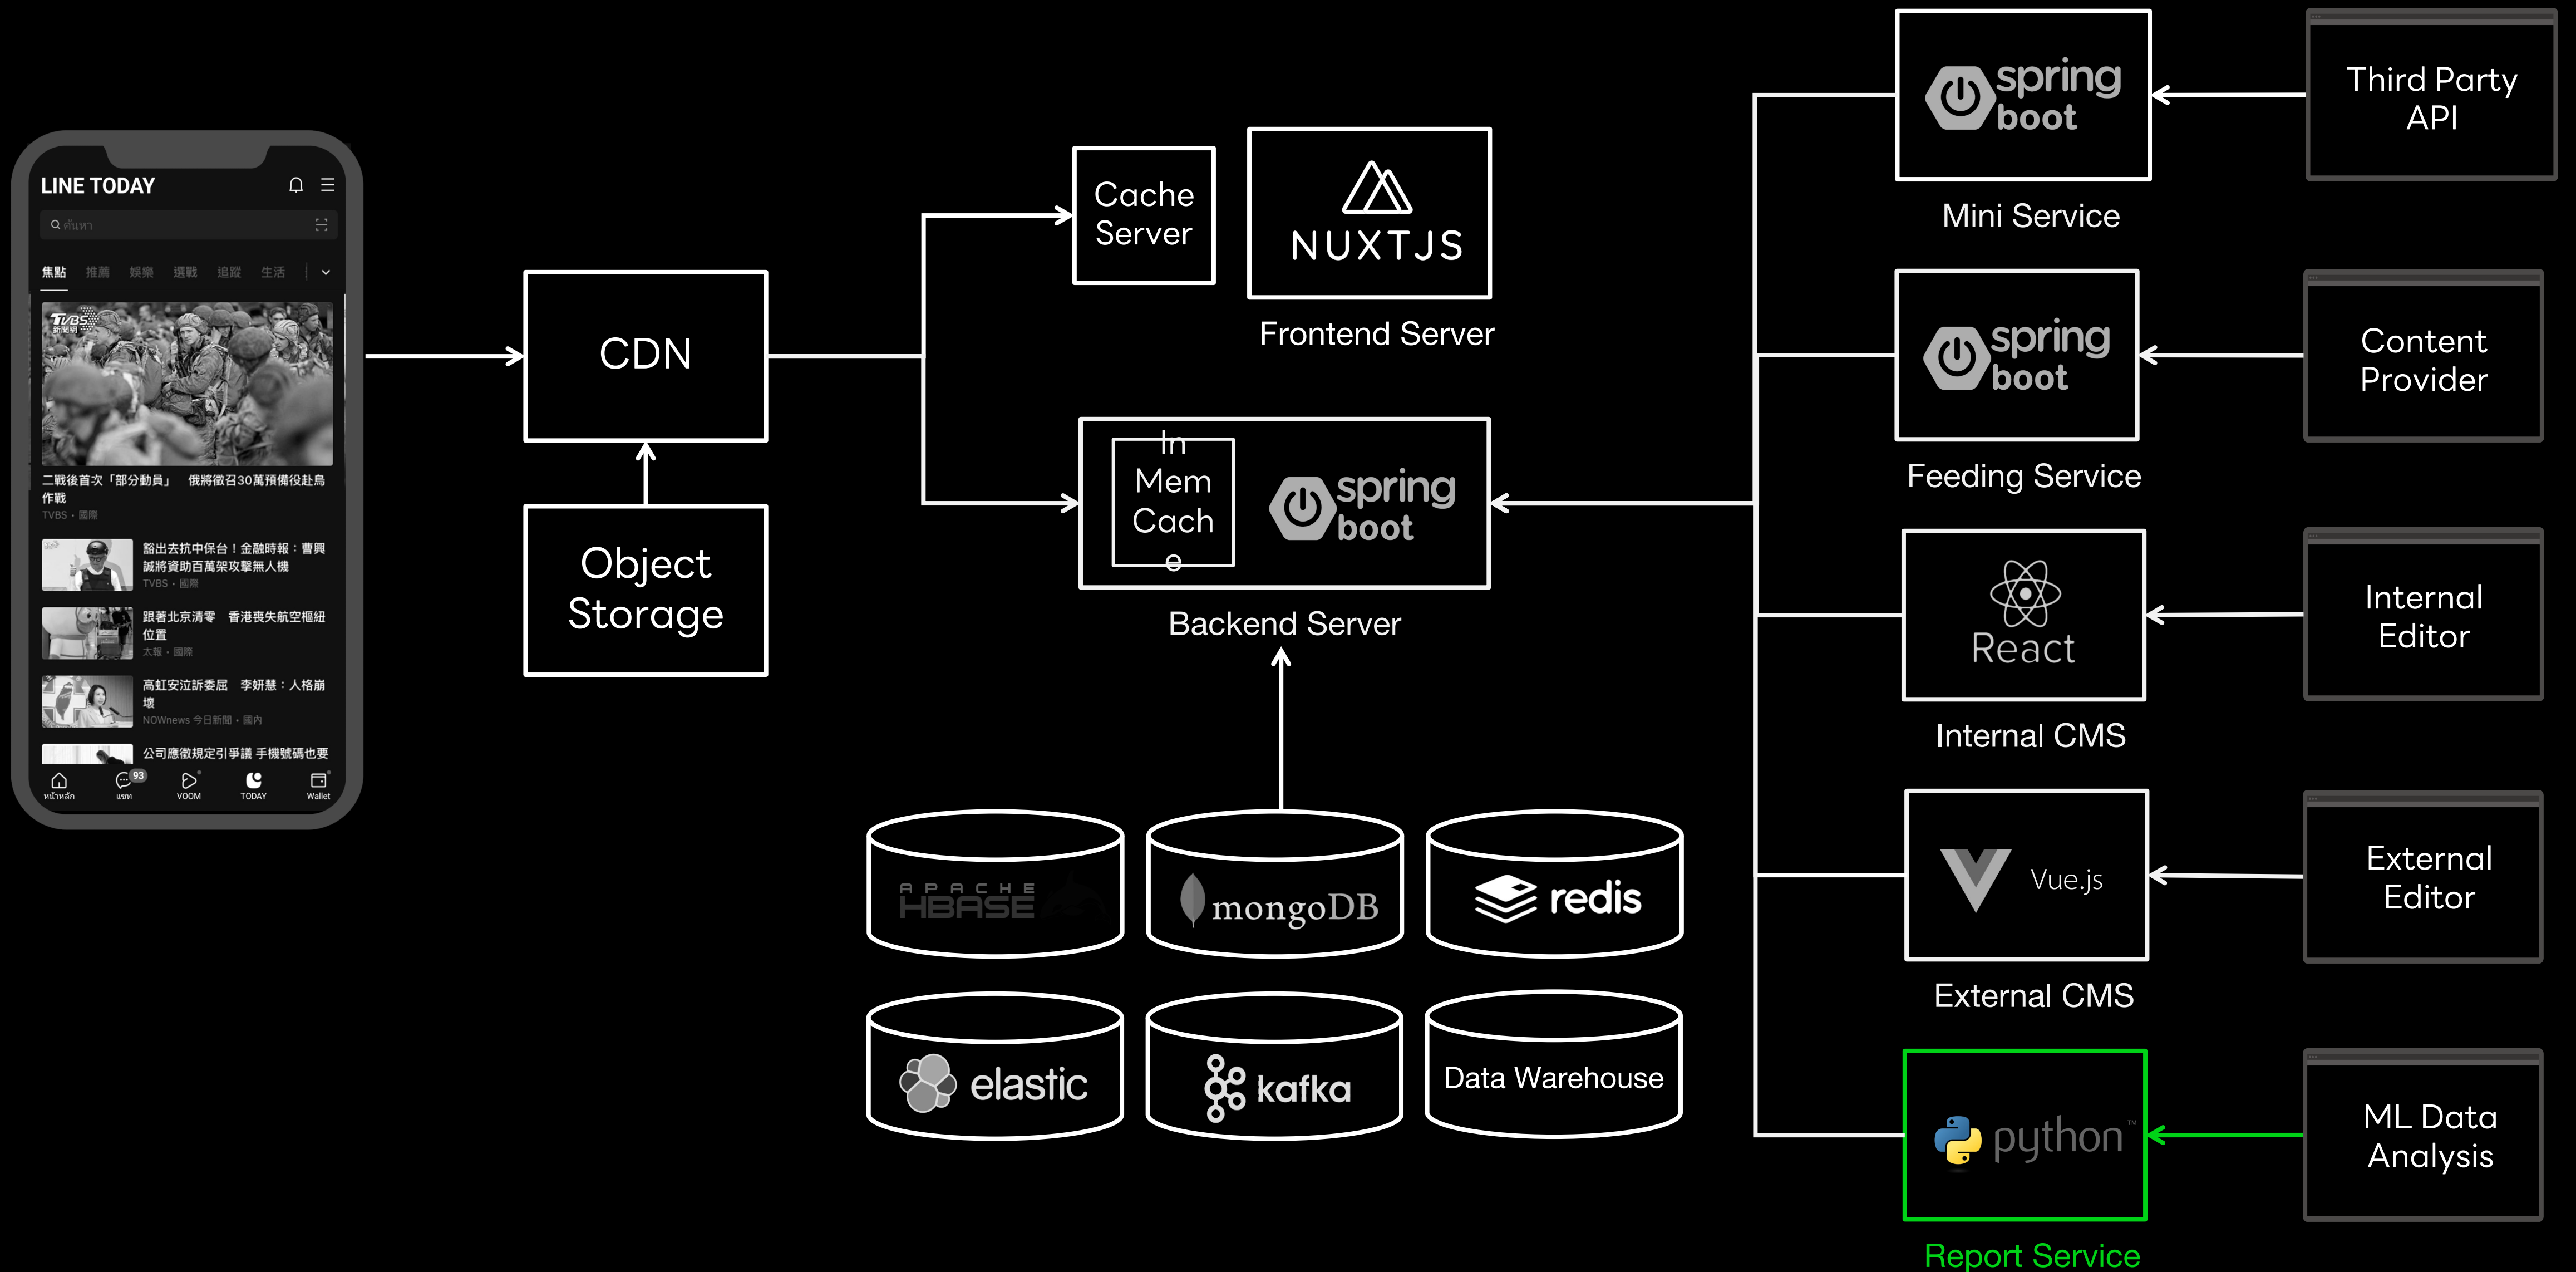
# LINE TODAY Architecture



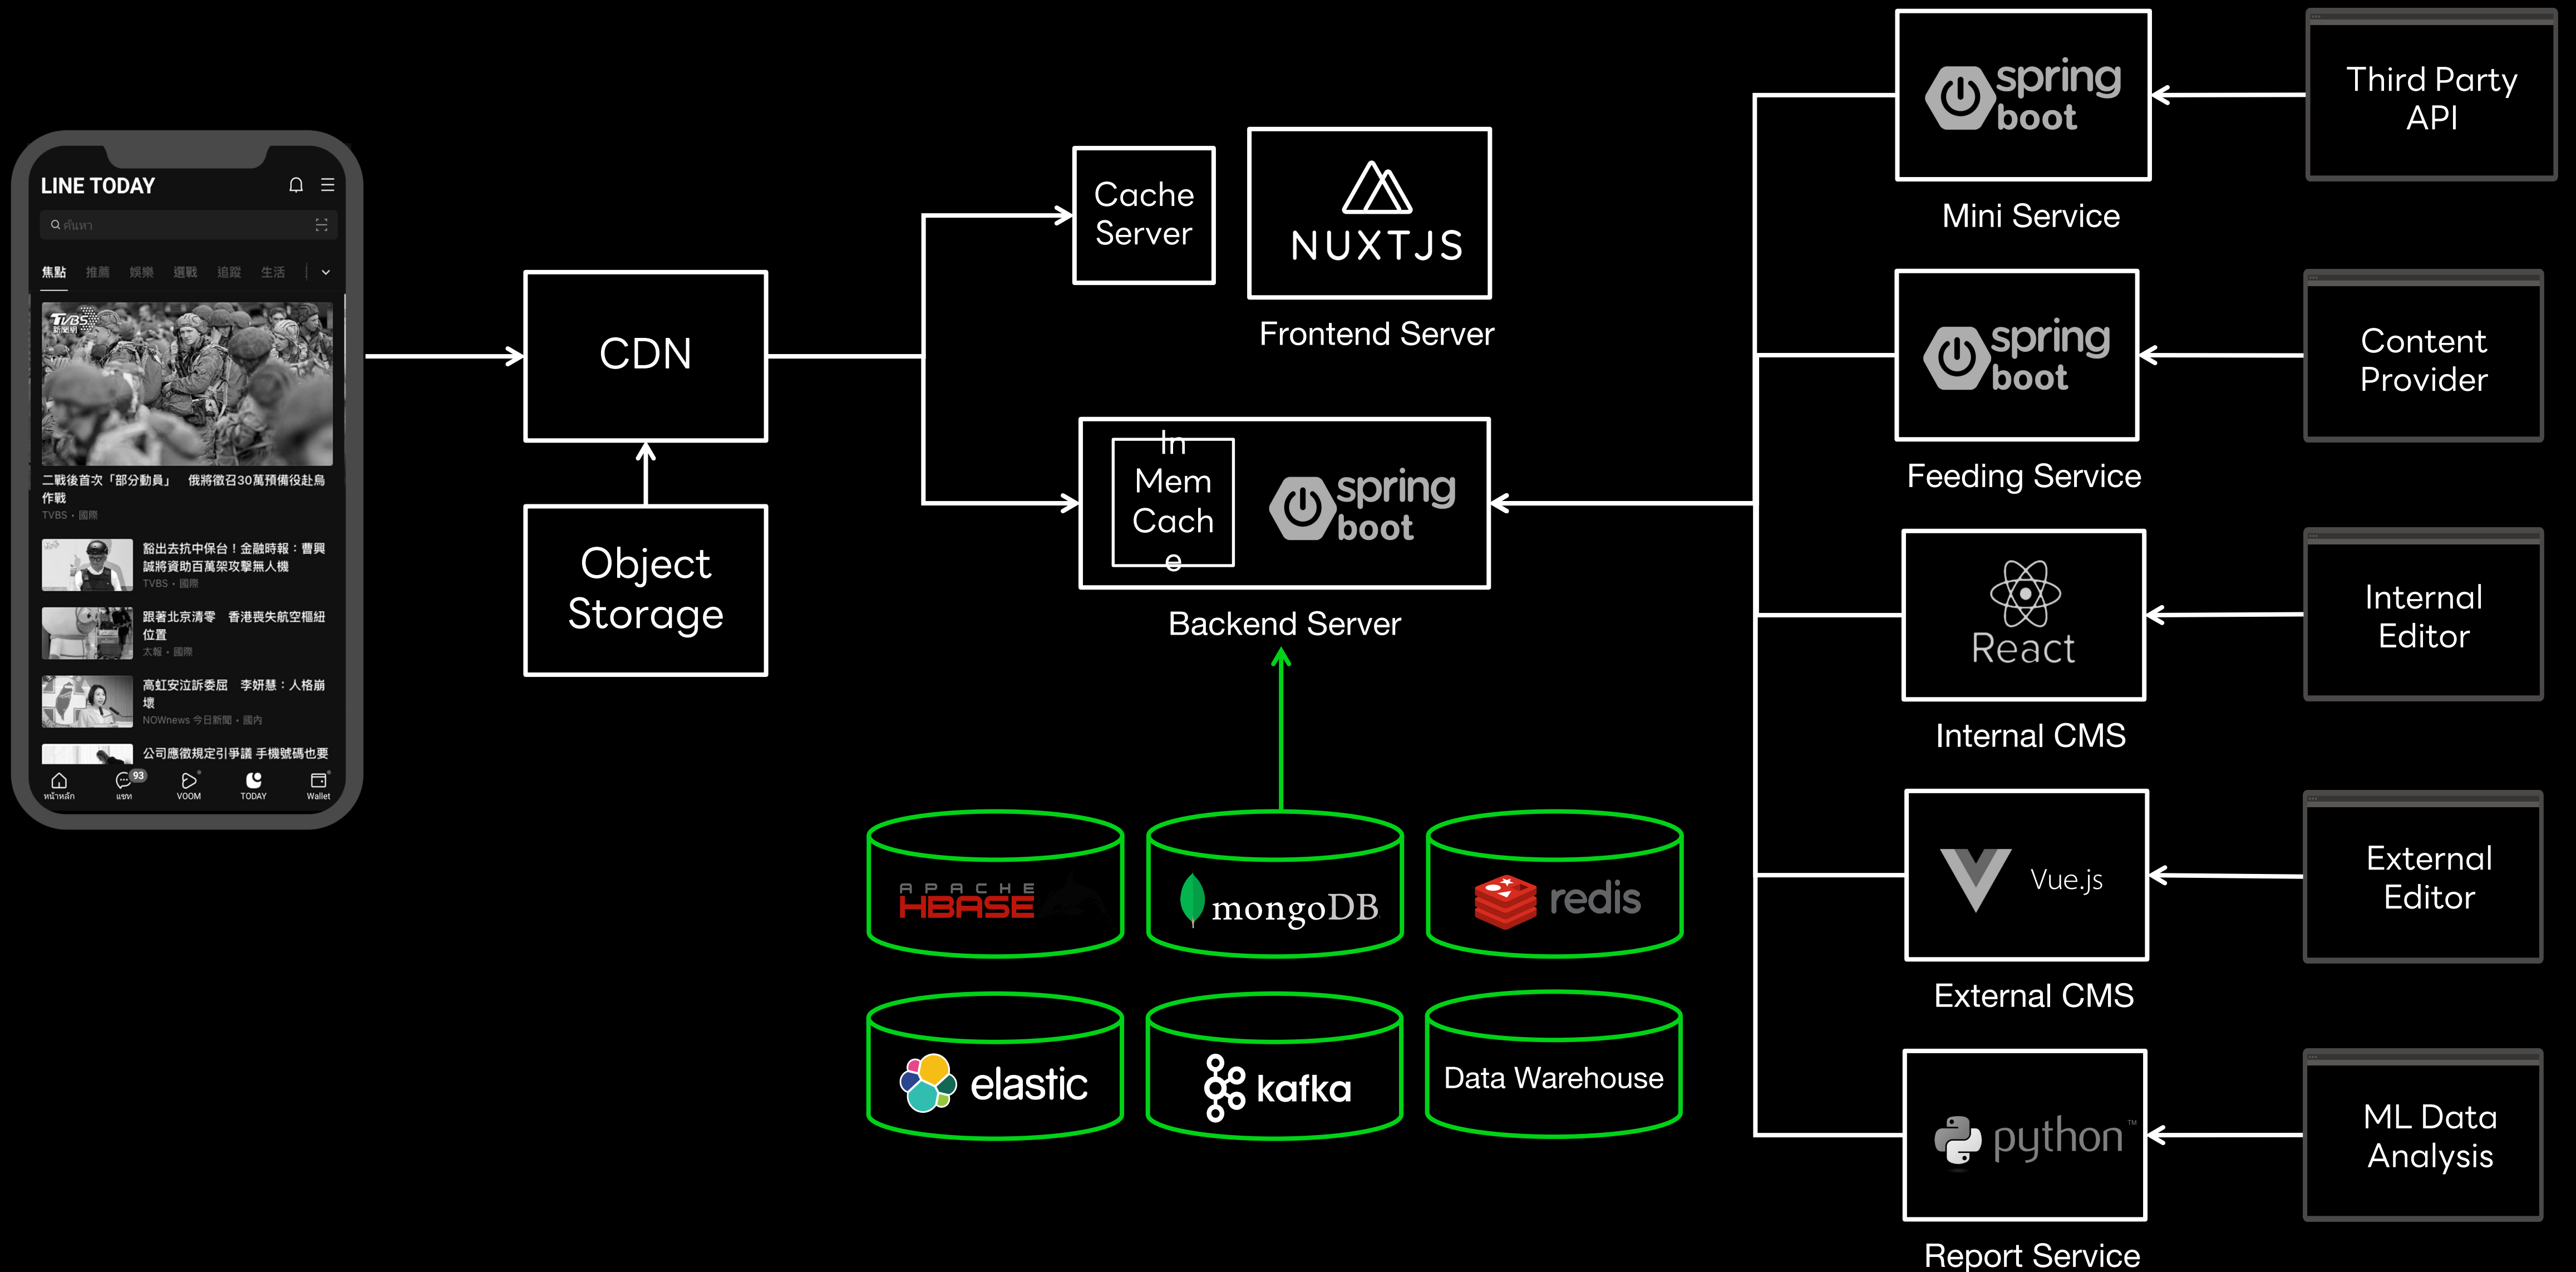
# LINE TODAY Architecture



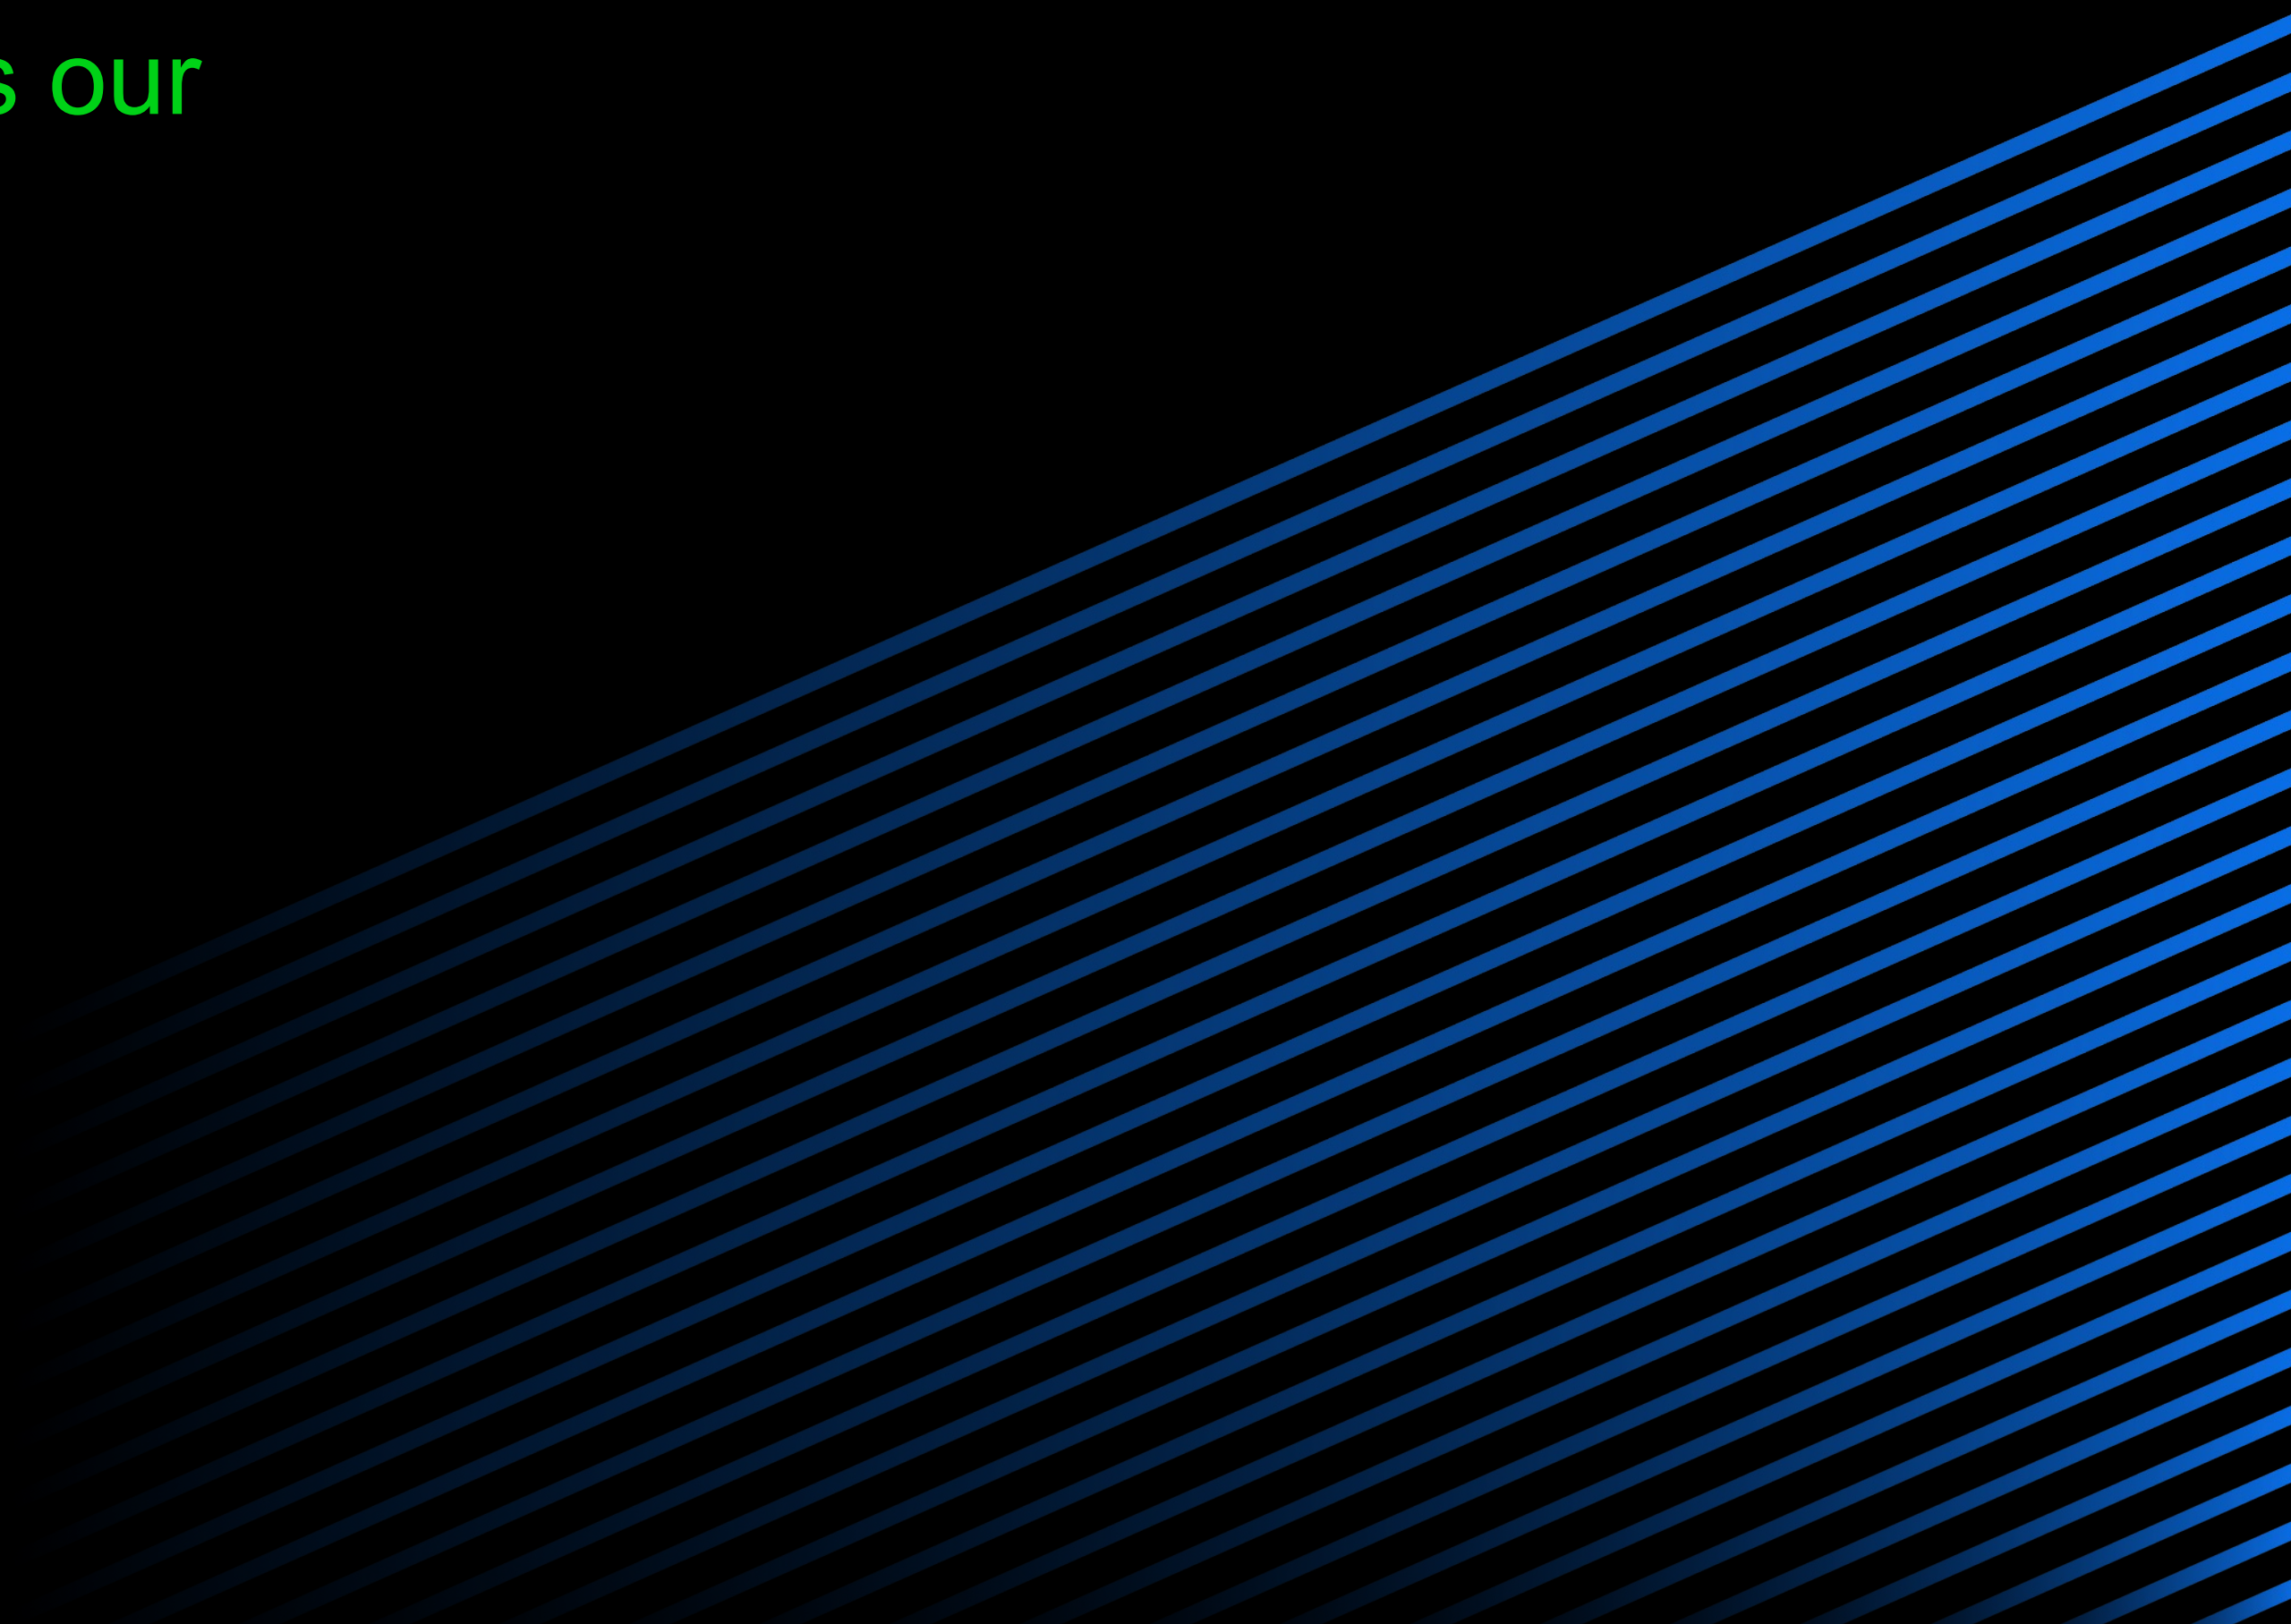
# LINE TODAY Architecture



# LINE TODAY Architecture



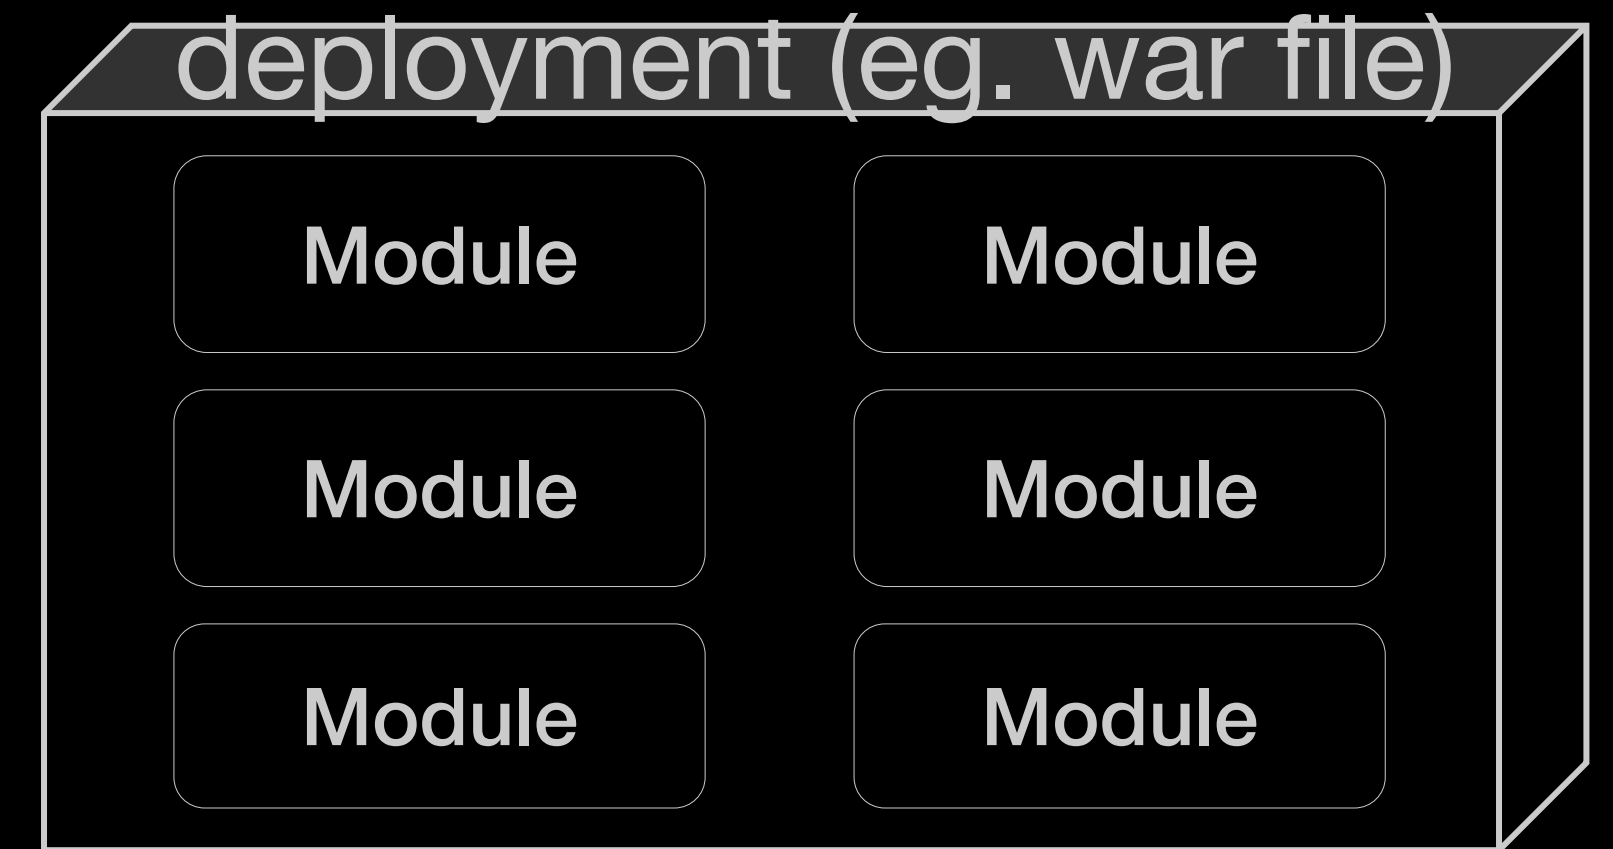
# Agenda

- LINE TODAY and its architecture
  - How miniservices and K8S changes our development and operation
    - Refactor to mini services
  - Lessons learnt
- 
- A decorative graphic consisting of numerous parallel blue diagonal lines that sweep across the bottom right portion of the slide, creating a sense of motion and depth.

# How to improve development and deployment efficiency?

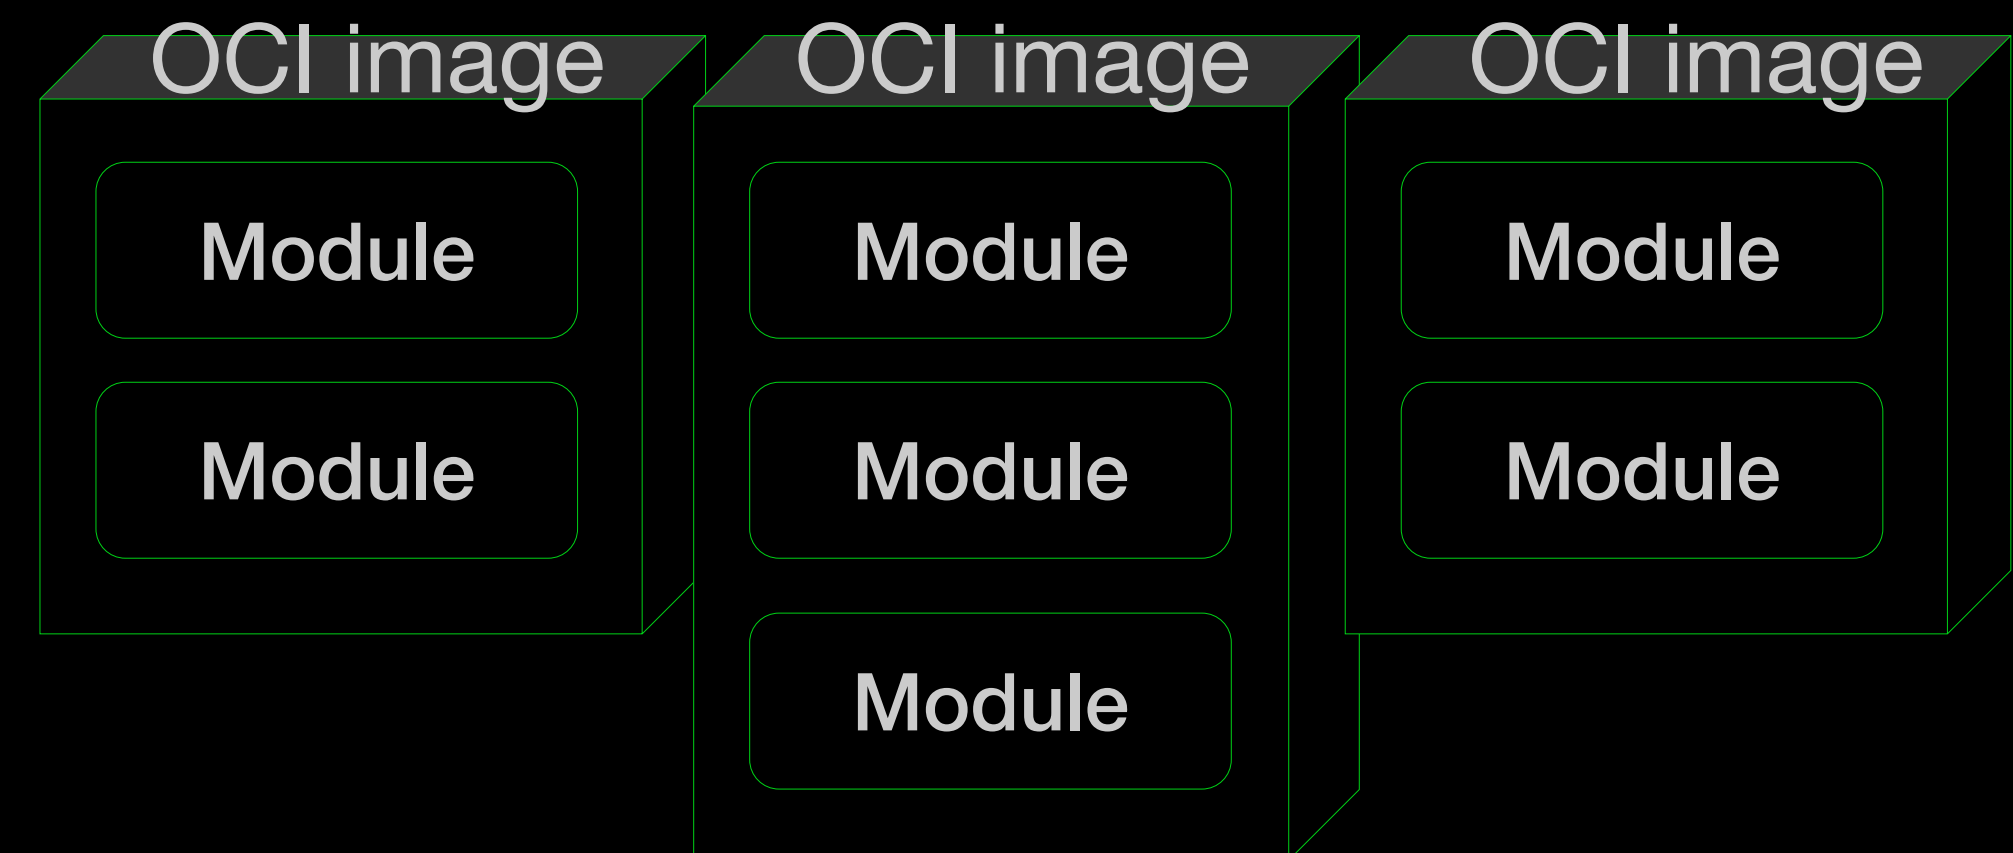
## Problem

Small change requires entire system rebuild and deployment

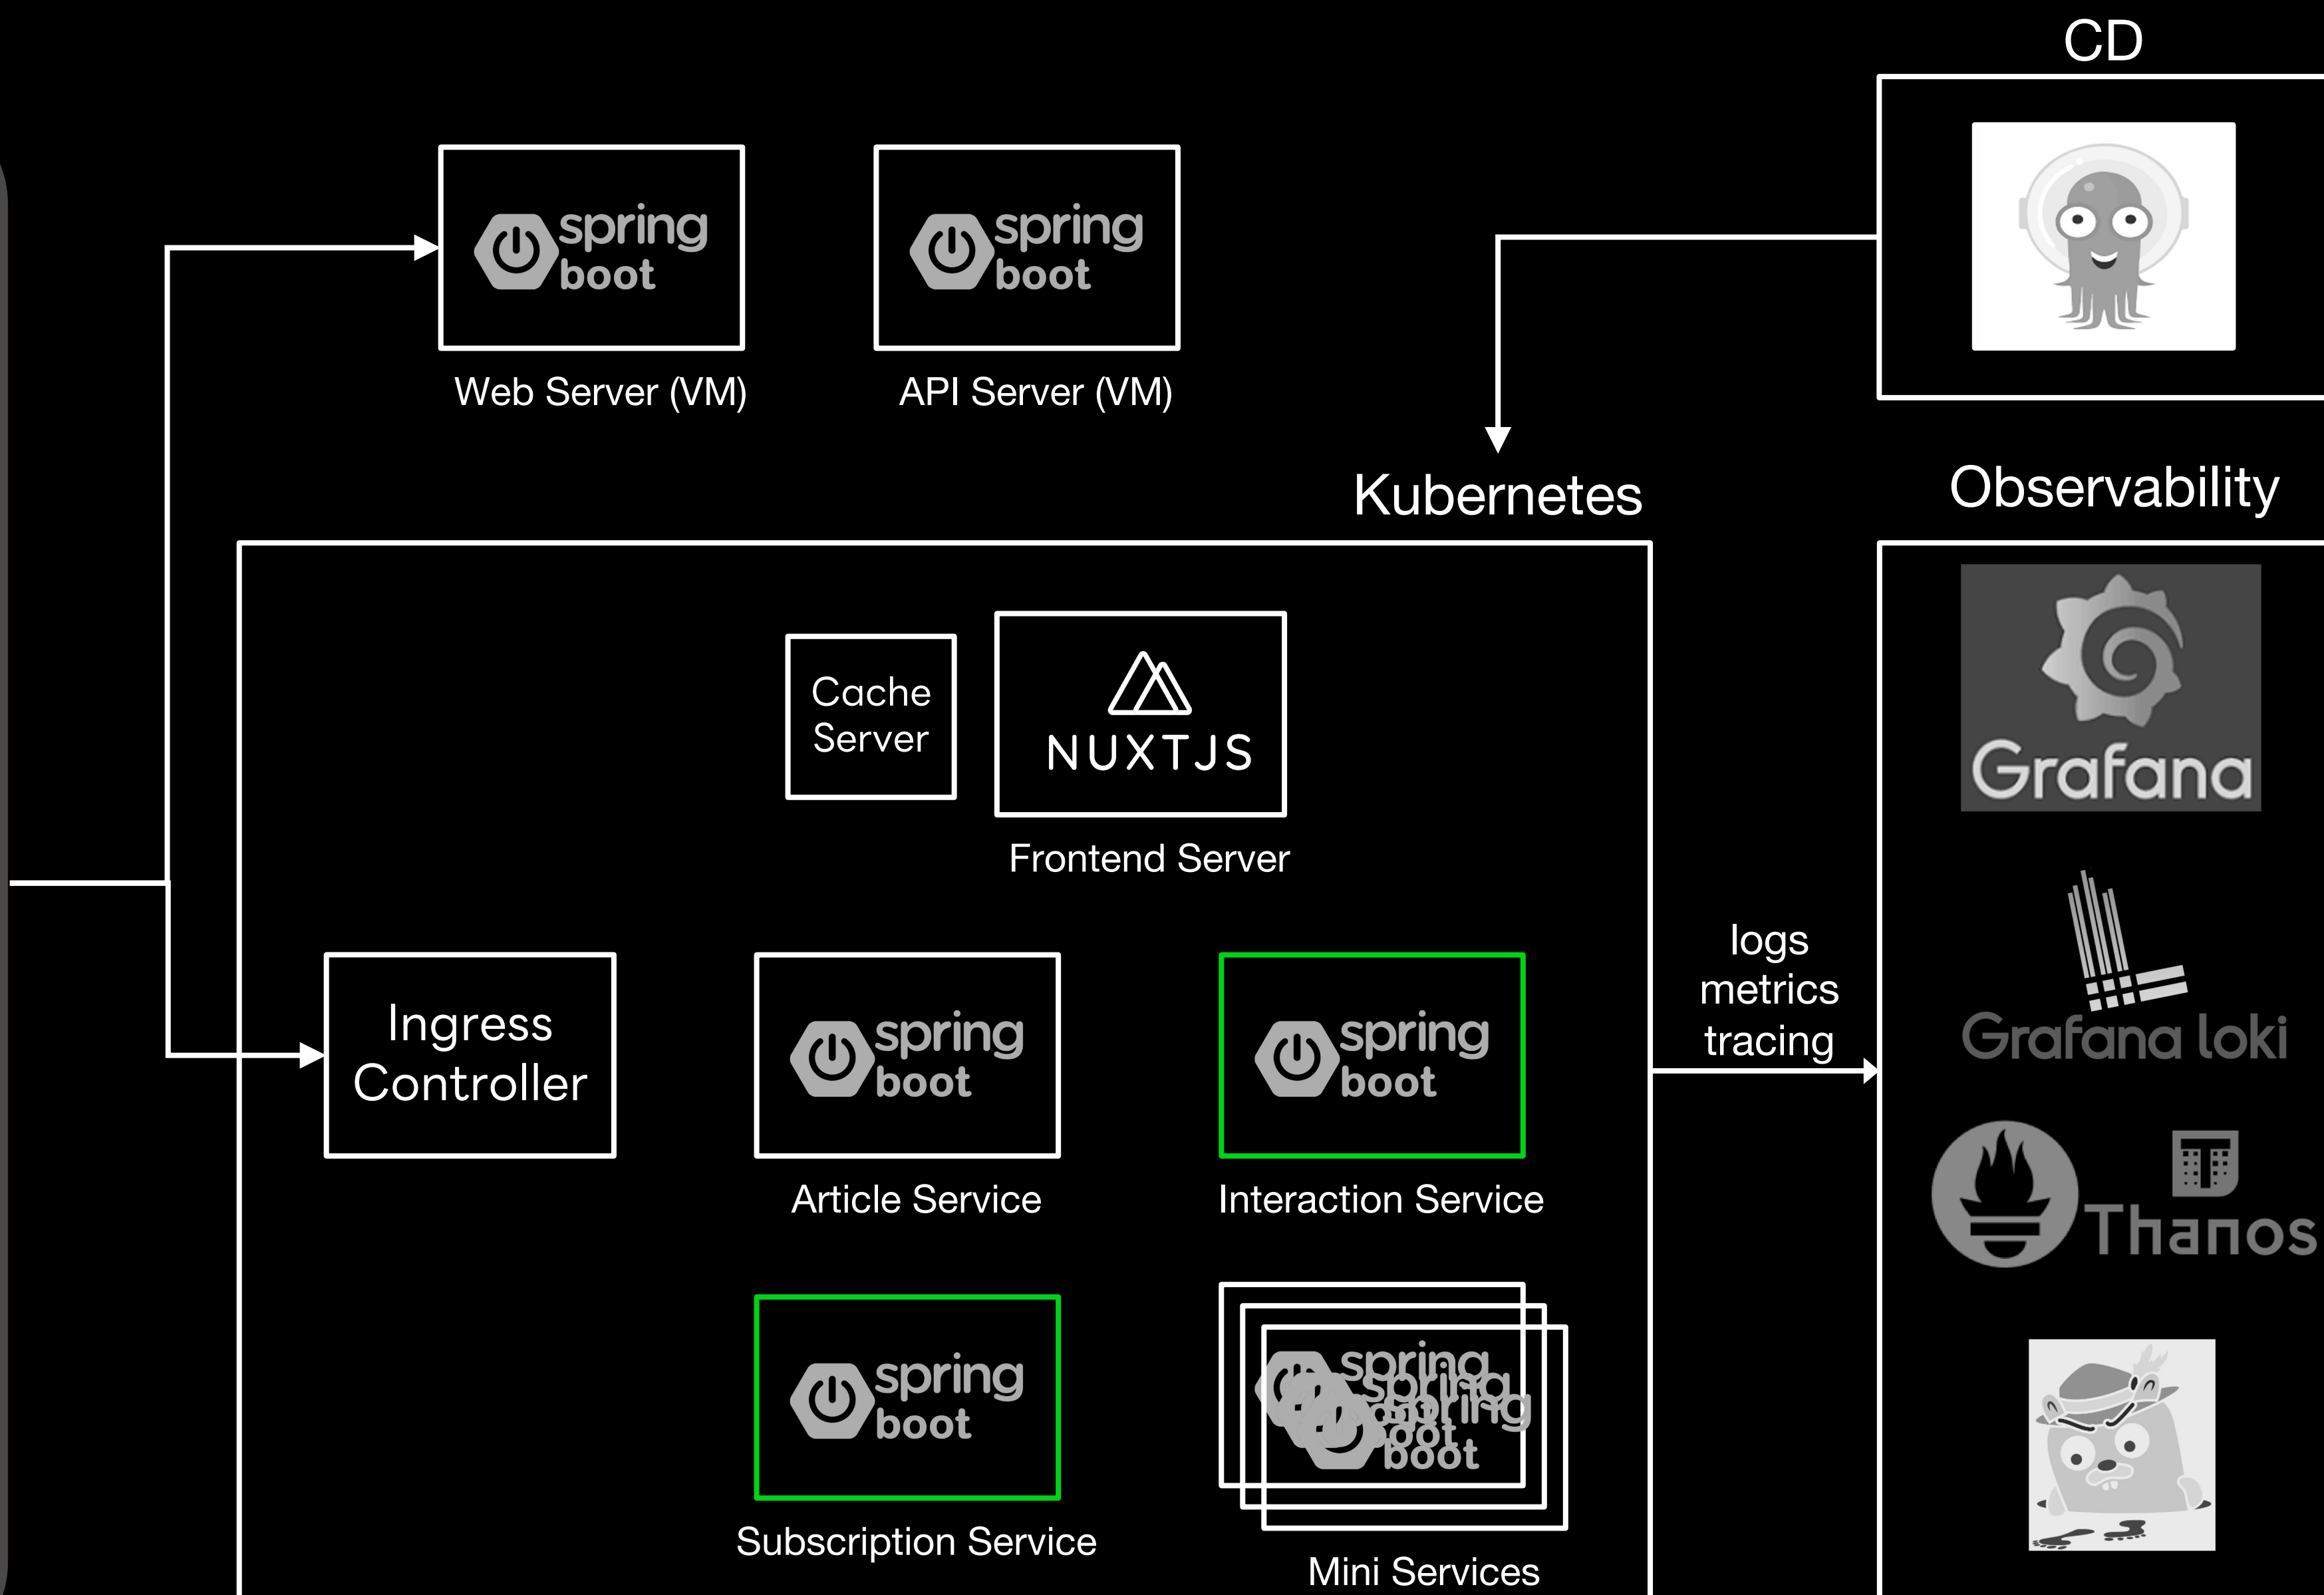


## Solutions

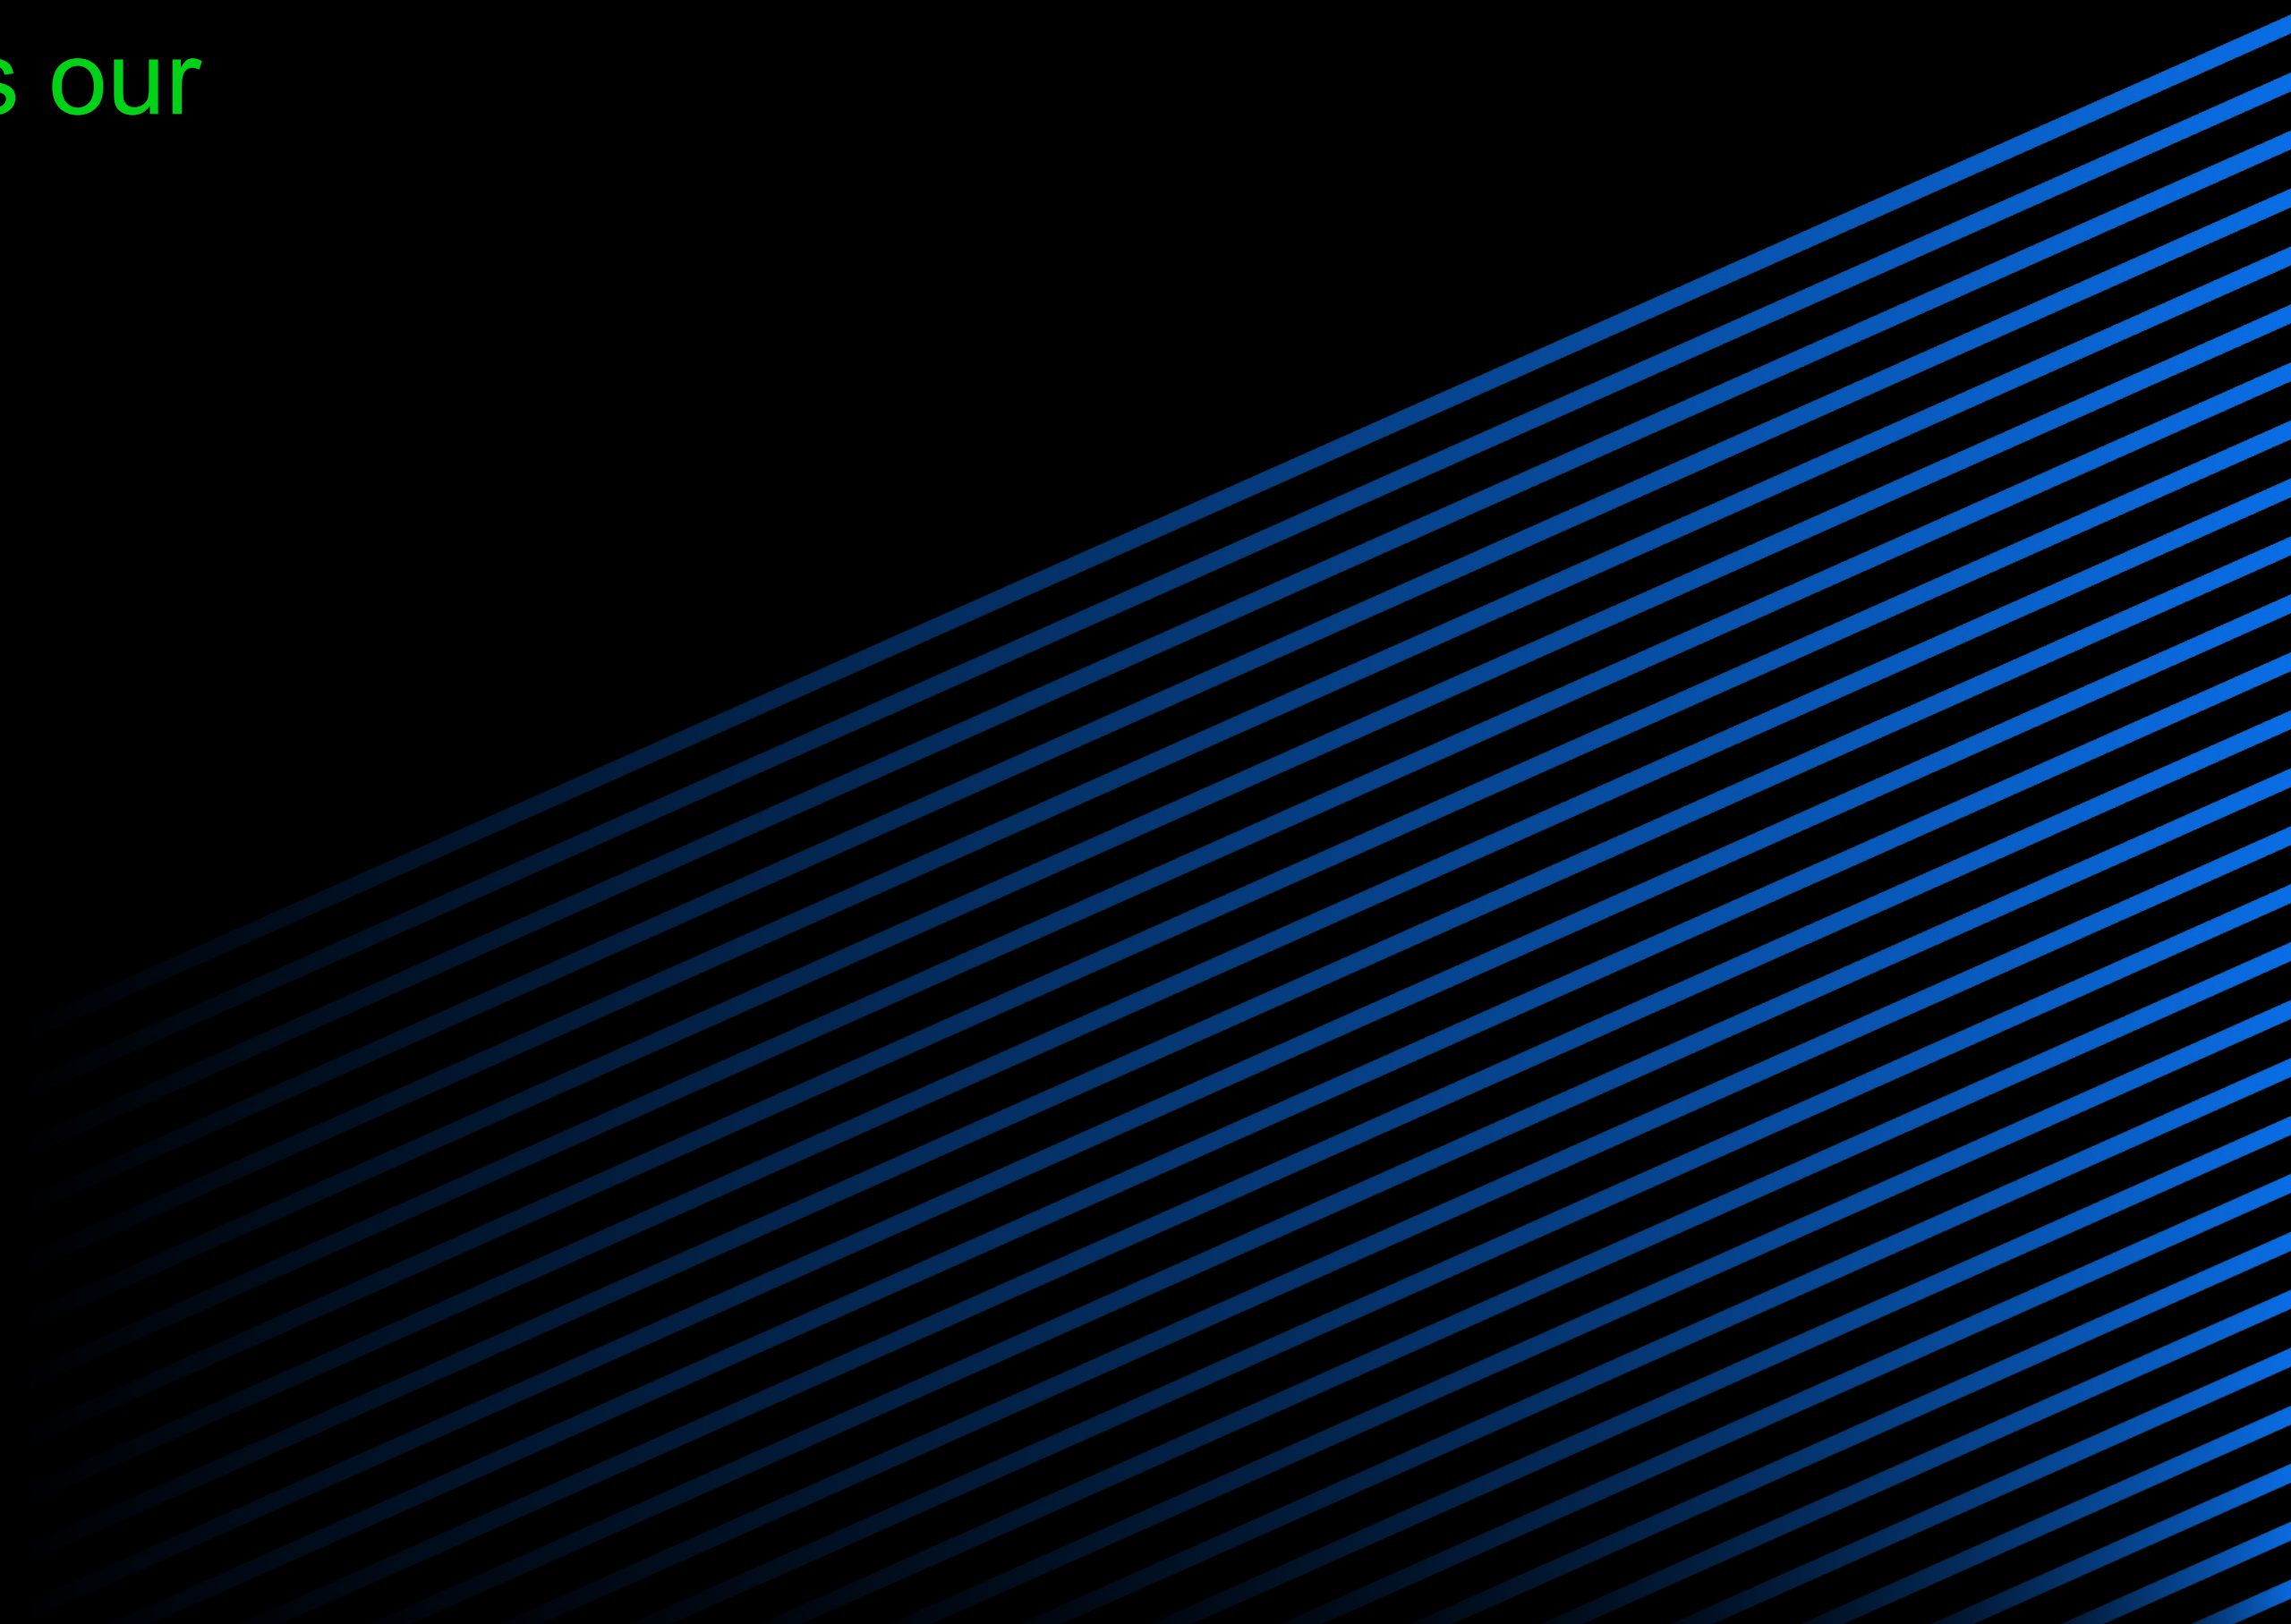
- Break down coarse-grained deployments into functionally cohesive mini services
- Move to Kubernetes



# Migrate to mini services and Kubernetes



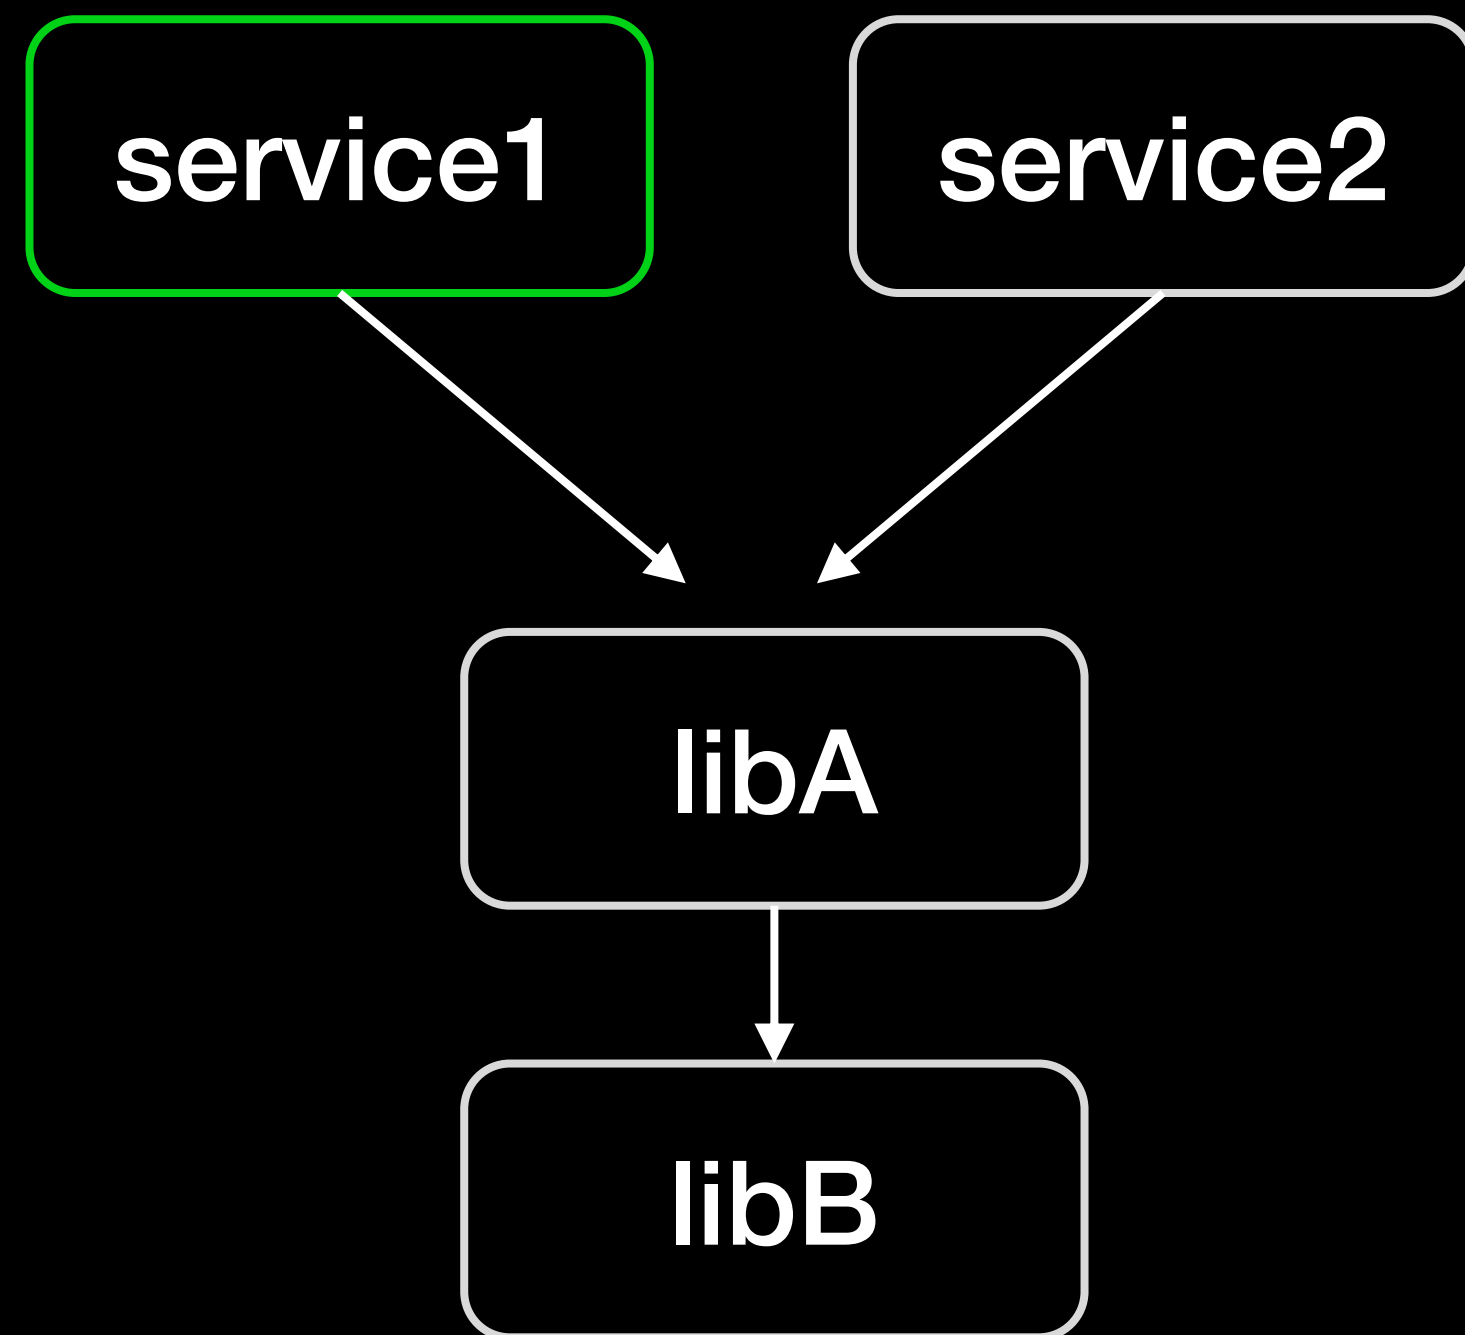
# Agenda

- LINE TODAY and its architecture
  - How miniservices and K8S changes our development and operation
    - Refactor to mini services
    - Refine CI/CD
  - Lessons learnt
- 
- A decorative graphic consisting of numerous parallel blue lines that fan out from the bottom right corner towards the top right, creating a sense of motion and depth against the dark background.

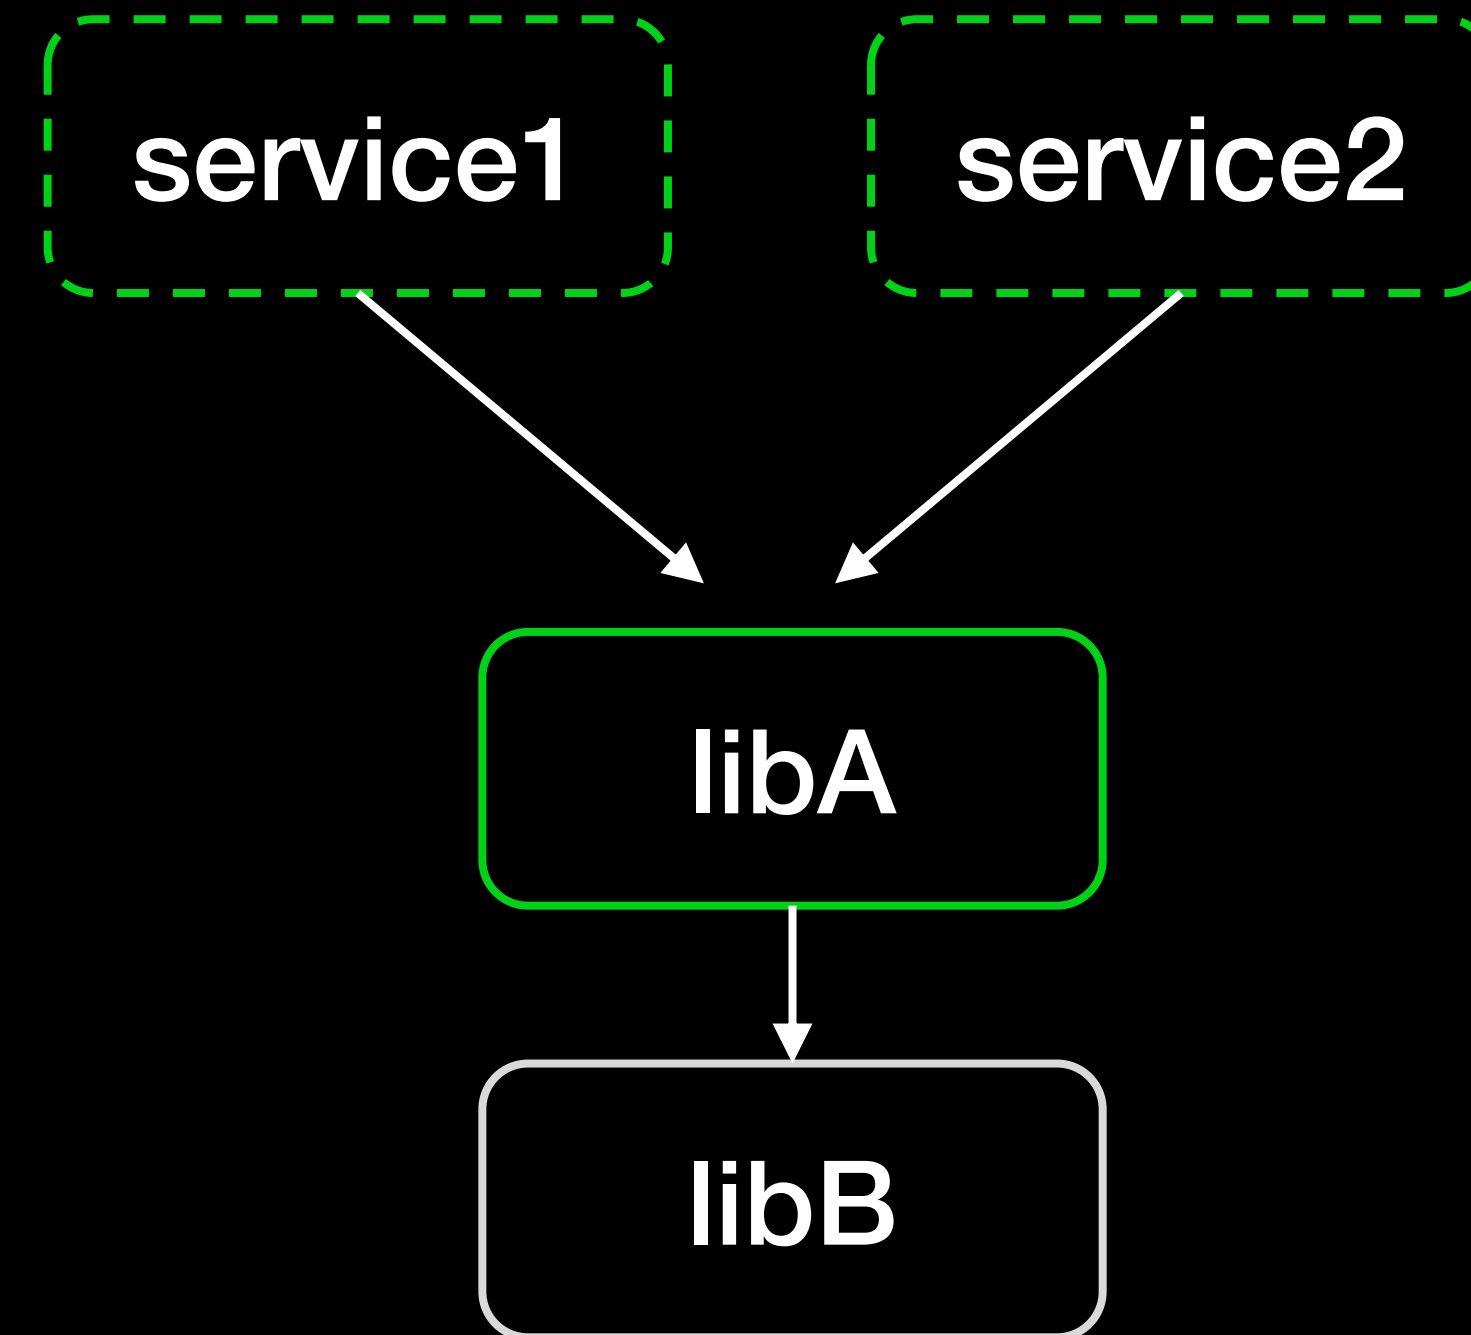


# Build and test what was changed

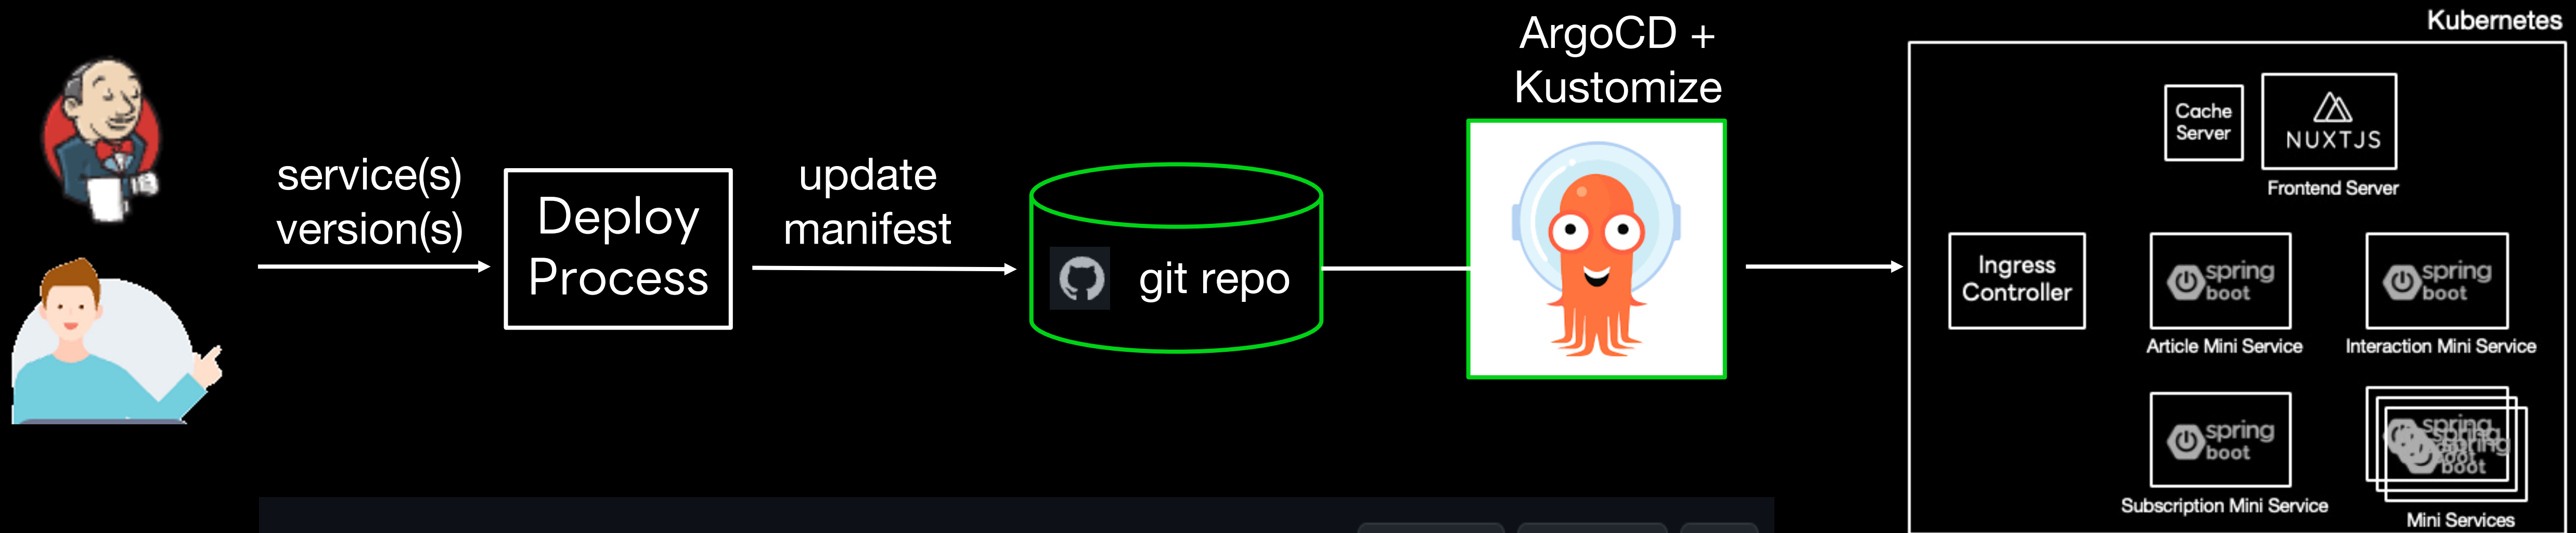
Changing service1  
=> rebuild service1



Changing libA => rebuild  
libA, service1, service2



# Manage deployment via GitOps

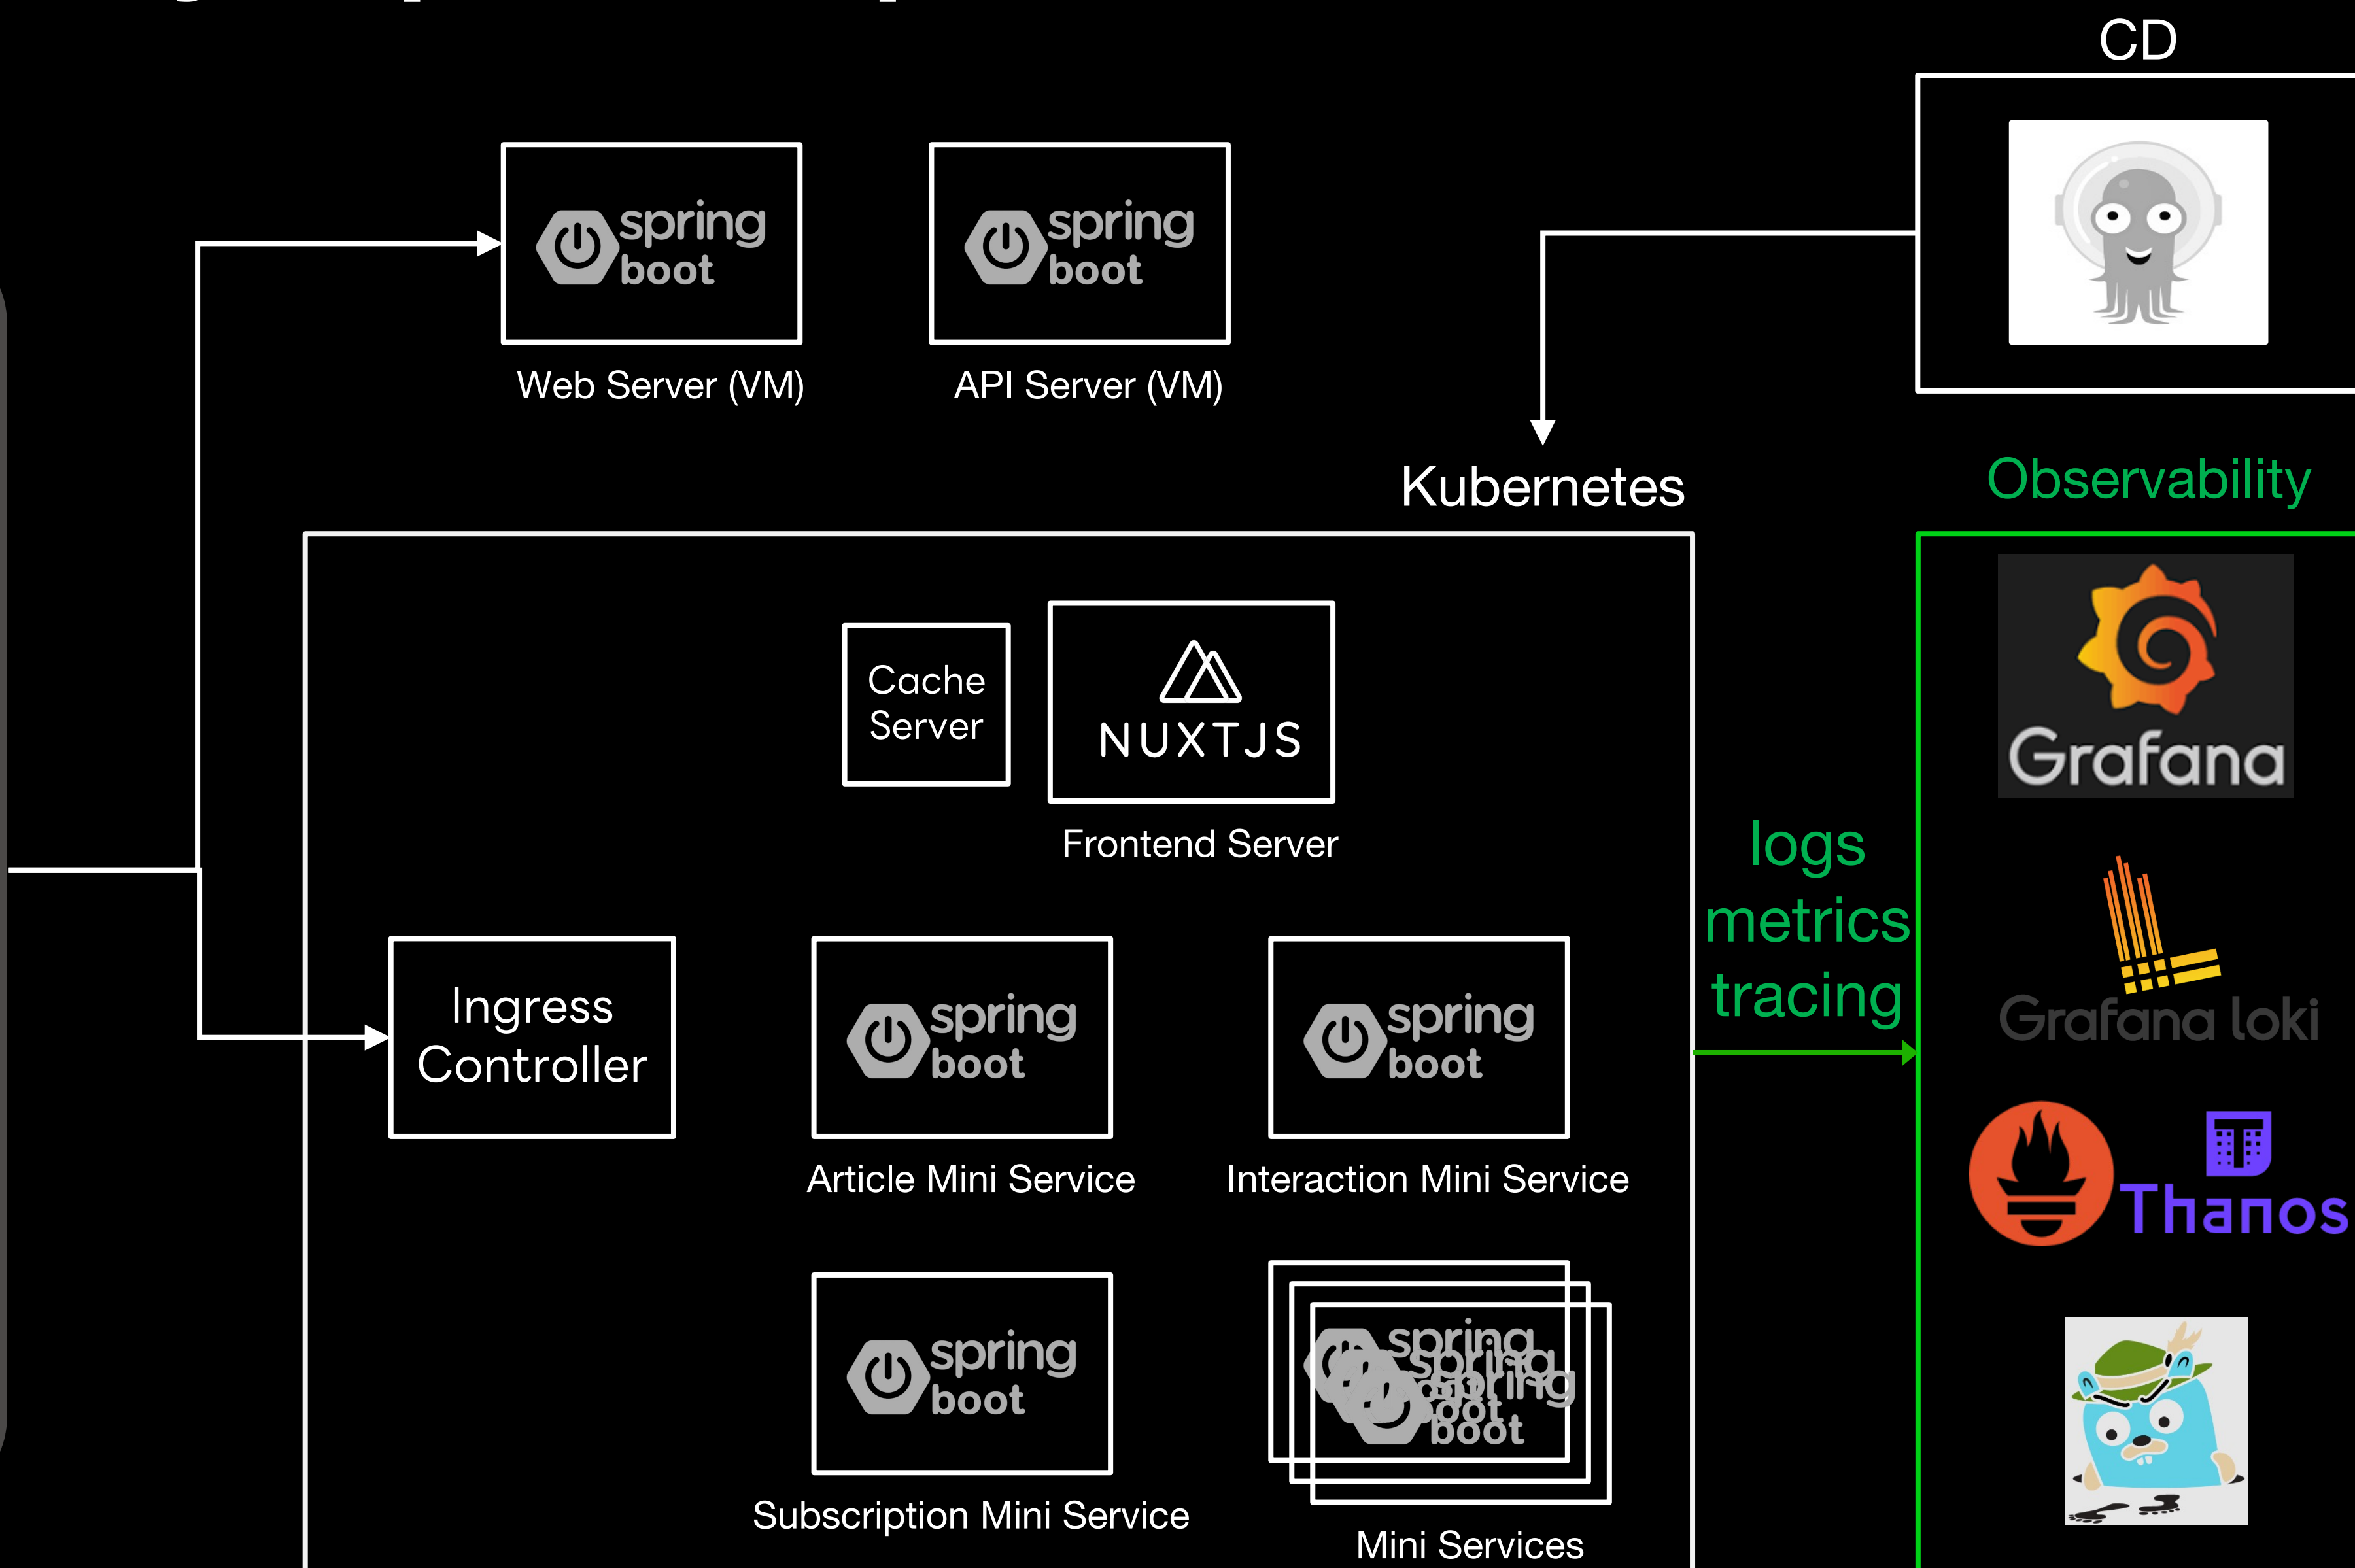


| Commit Hash                          | Message  | Time        |
|--------------------------------------|--|-------------|
| 93feac88ab7c1f0958f1322ccc0940d1e559 | 4667d3a  | 3 hours ago |
| ..                                   |  |             |
| alpha                                | [alpha] 4c5c93feac88ab7c1f0958f1322ccc0940d1e559   | 3 hours ago |
| beta                                 | [beta] 4c5c93feac88ab7c1f0958f1322ccc0940d1e559    | 3 hours ago |
| rc                                   | [rc] 00ad13169c8932ef8b683d62d67f7714cdff5865      | 4 days ago  |
| release                              | [release] fabf018ecdff20ecfda5b939dc37ff1f4b94f52e | 5 days ago  |

# Agenda


- LINE TODAY and its architecture
- How miniservices and K8S changes our development and operation
  - Refactor to mini services
  - Refine CI/CD
  - Integrate with observability
- Lessons learnt

# Observability improves operation efficiencies



# Observability - service RPS / latency metrics

### RPS



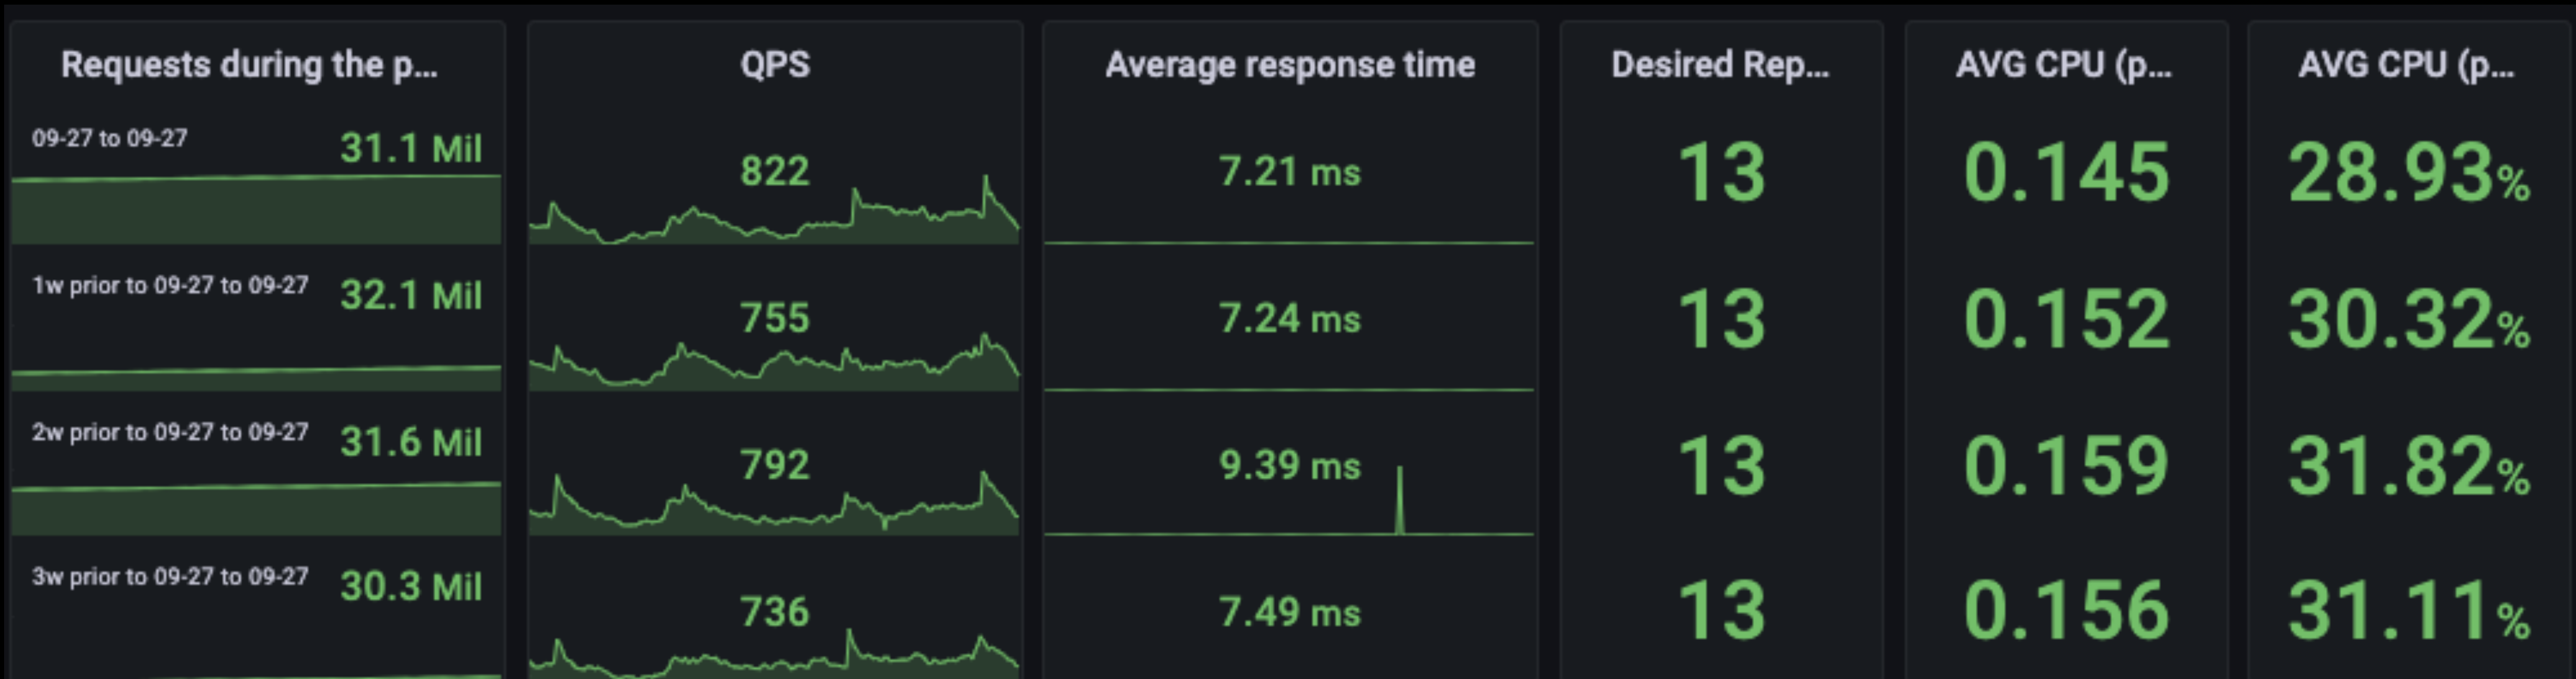
|     | Mean   | Max    | Last * |
|-----|--------|--------|--------|
| 2xx | 5.79 K | 7.77 K | 5.70 K |
| 3xx | 2.50   | 6.92   | 2.29   |
| 4xx | 13.1   | 20.1   | 13.5   |
| 5xx | 0.0308 | 0.111  | 0      |

### Latency



|                 | Mean    | Max     |
|-----------------|---------|---------|
| 95th Percentile | 97.8 ms | 101 ms  |
| 50th Percentile | 51.4 ms | 52.6 ms |
| Average         | 16.2 ms | 45.7 ms |

# Observability - metrics week-over-week

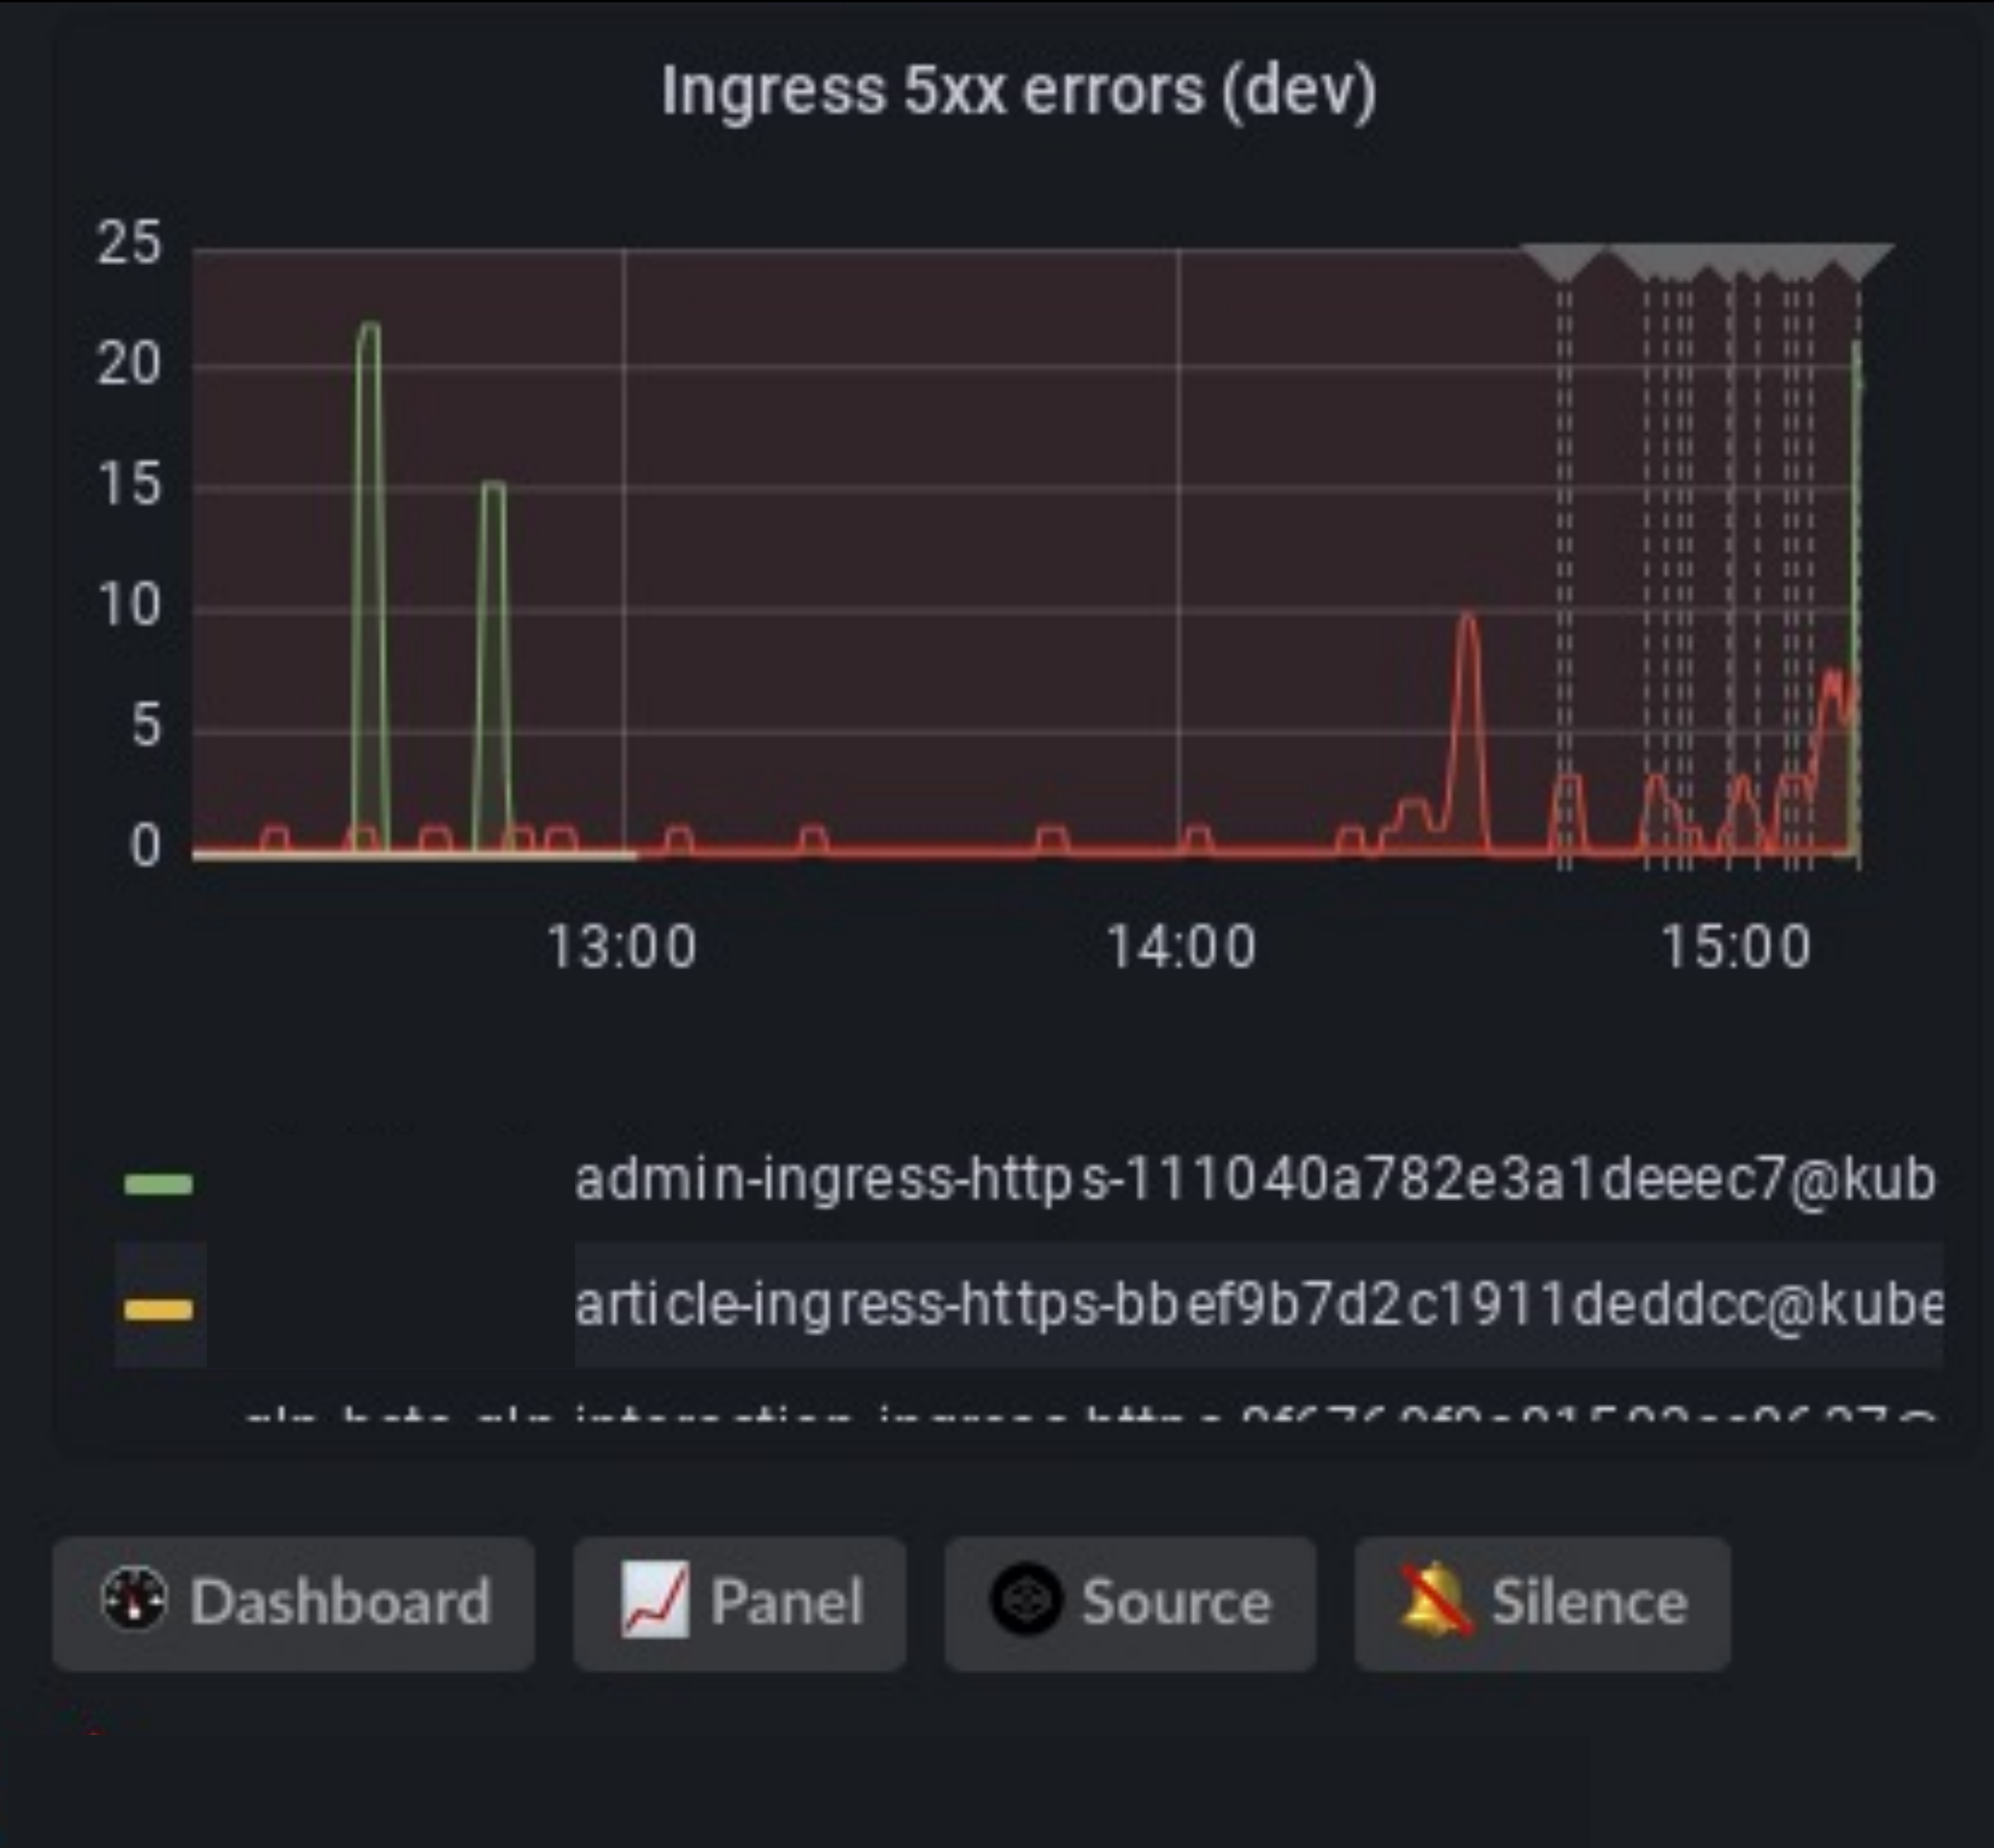


# Observability - abnormal spikes

rate of requests returning 500

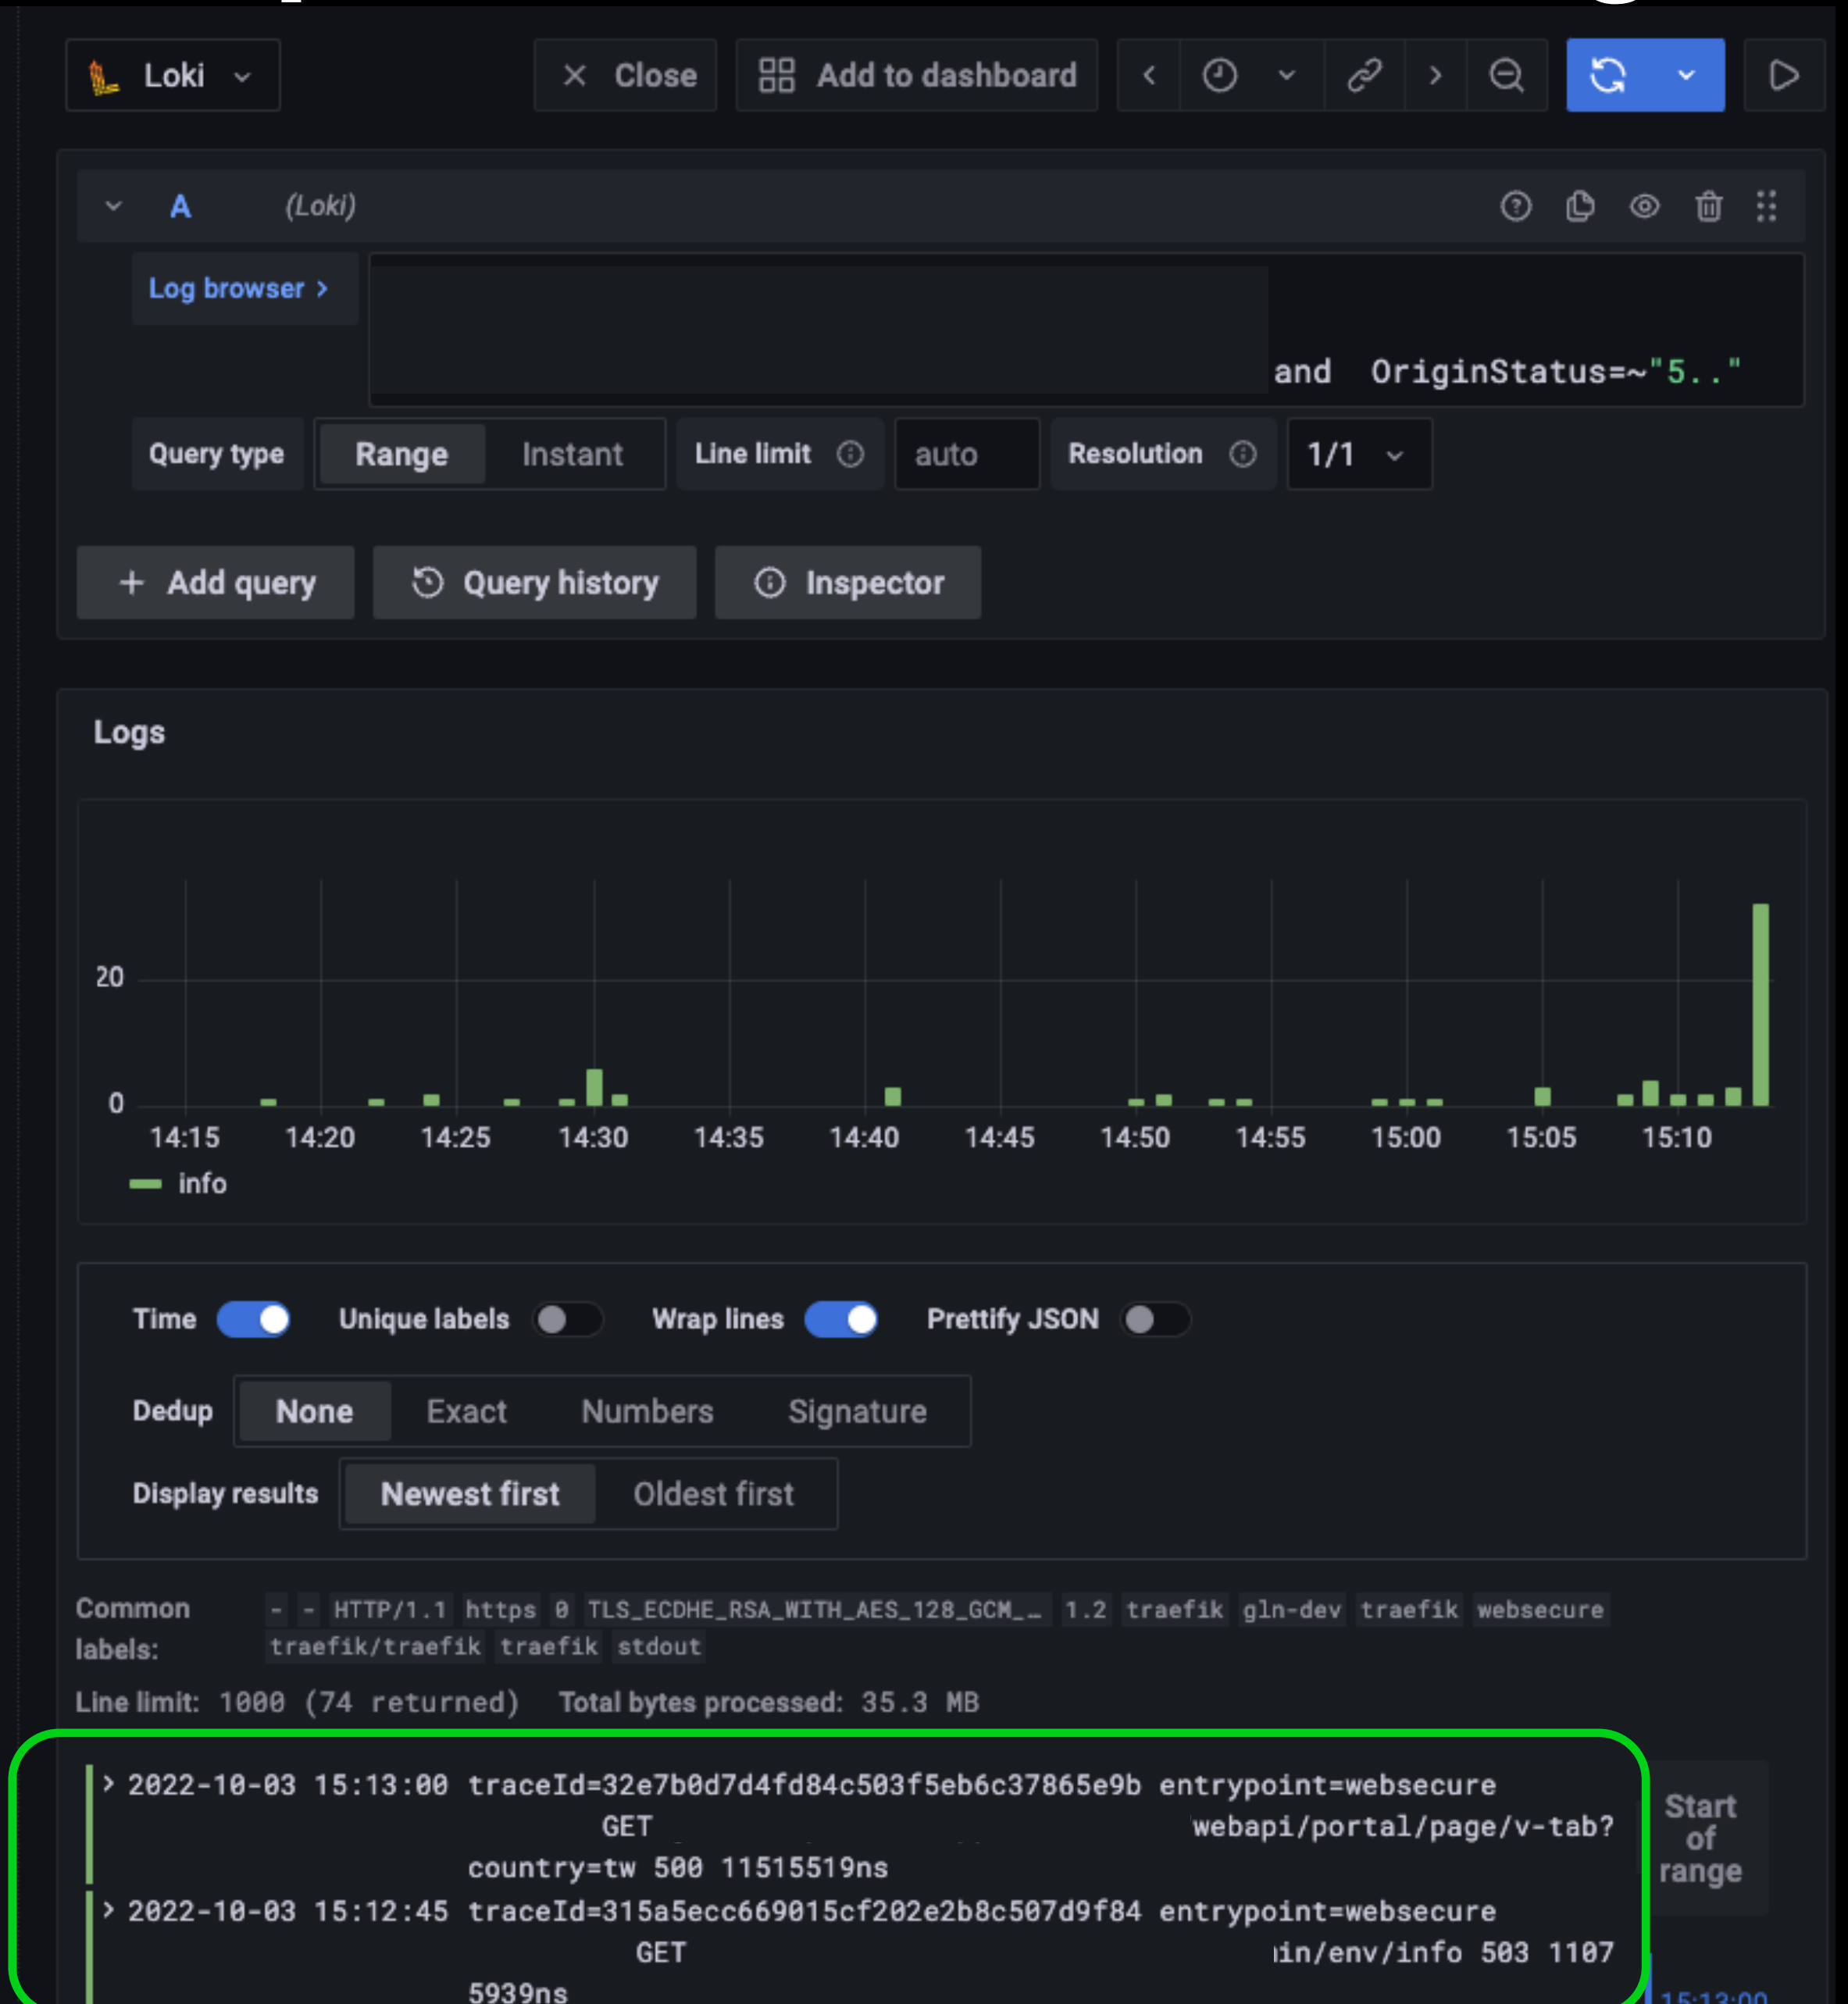
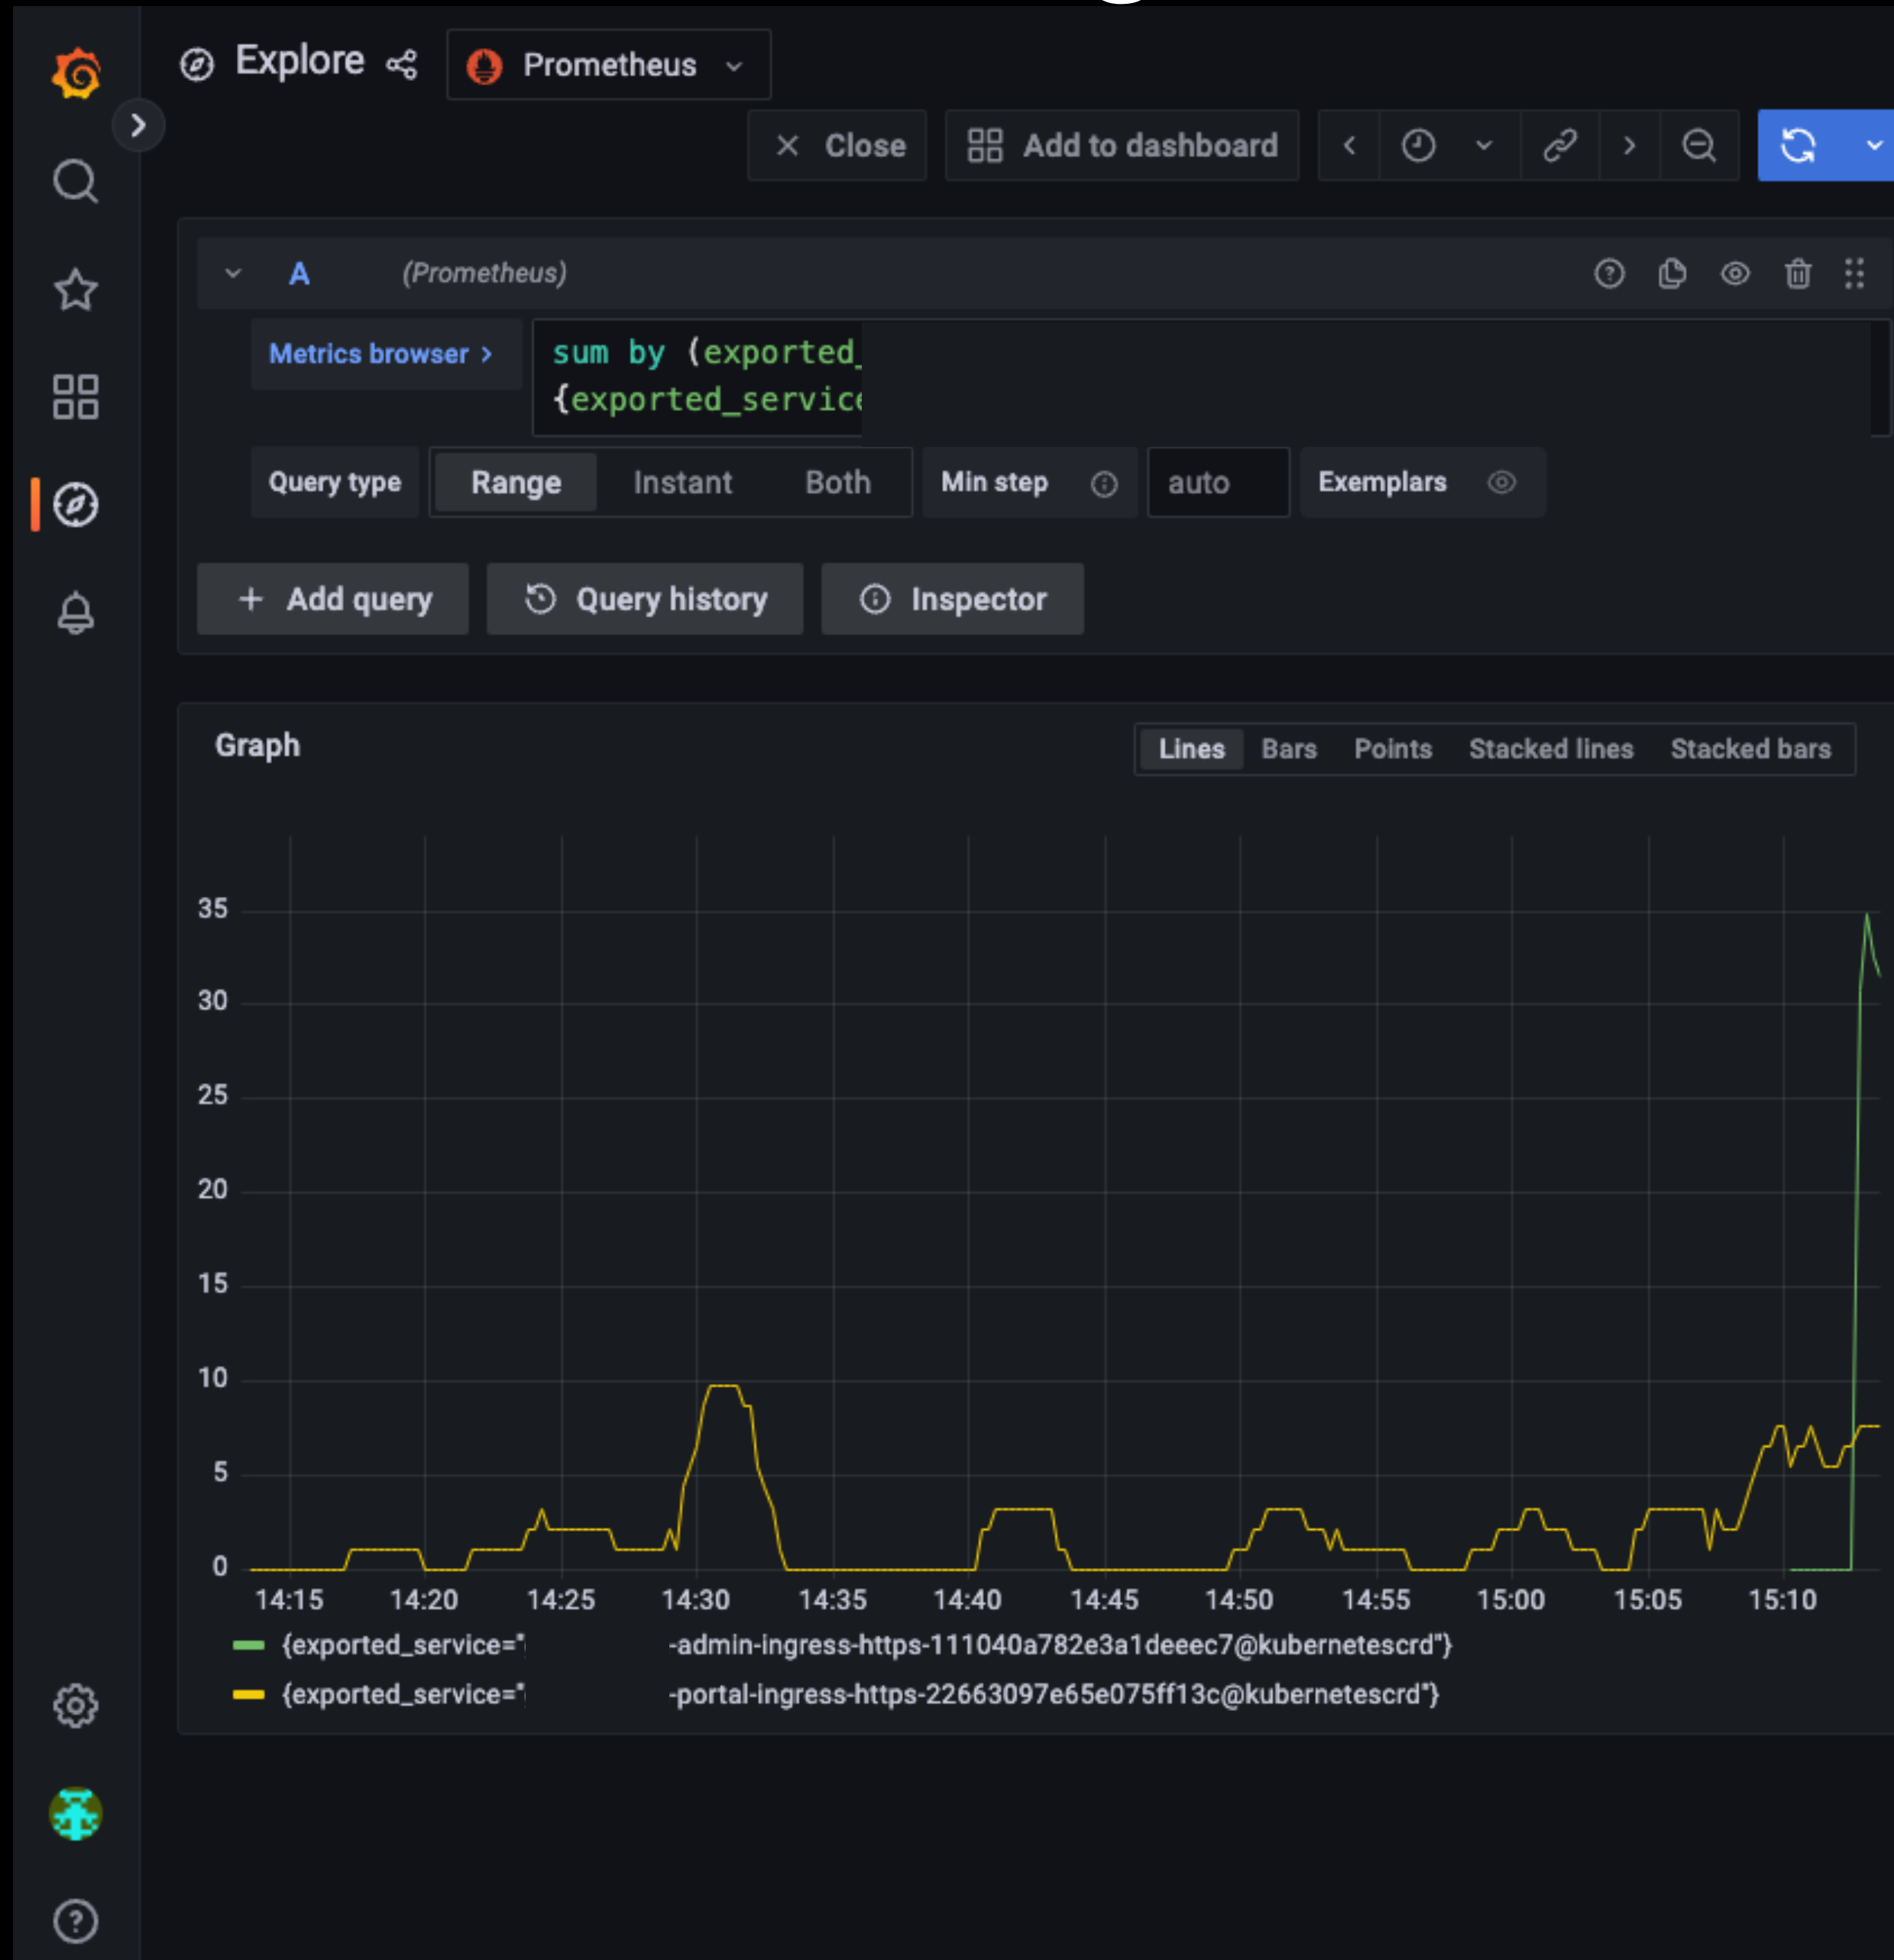


# Troubleshooting - 1. receive alert from slack





# Troubleshooting - 2. link to alert panel and access logs



# Troubleshooting - 3. open trace viewer

2022-10-03T07:12:31Z

Detected fields

- TraceID: 289199fb067d6cff70b9520b5db62017
- Tempo
- Time
- entrypoint
- traceId
- tsNs

traefik

- traefik forward
- gln-portal
- gln-portal IveEndController

r.getIveSetting Service: gln-portal Duration: 5.09ms Start Time: 62.62ms (14:21:24.621) Child Count: 10

Logs for this span

Attributes: error = true | internal.span.format = proto | otel.library.name = io.opentelemetry.spring-webmvc-3.1 | otel.library.version = 1.18.0-alpha...

Resource: cluster = | container.id = 3375176f6b4e99f1406673b0bad76e99121bfe6868f152d88d830a17fbc2b47 | datasource = tw-gln | h...

Events (1)

- 67.7ms: message = exception | exception.stacktrace = java.lang.NullPointerException at |Servic...

Log timestamps are relative to the start time of the full trace.

SpanID: d3e7123f937b0a4e

gln-portal HGET (674.37μs)

- gln-portal HTTP GET (9.99ms)

  - gln-article /internal/a

    - gln-article Article

      - gln-article fin
      - gln-article fin
      - gln-article HC

- gln-portal GET (385.61μs)
- gln-portal find global\_new

674.37μs

- 9.99ms
- 8.26ms
- 7.54ms
- 1.79ms
- 1.01ms
- 574.57μs
- 385.61μs
- 1.91ms

# Troubleshooting via observability

Alerts

Received alerts from slack

Metrics

Check error source and time period

Logs

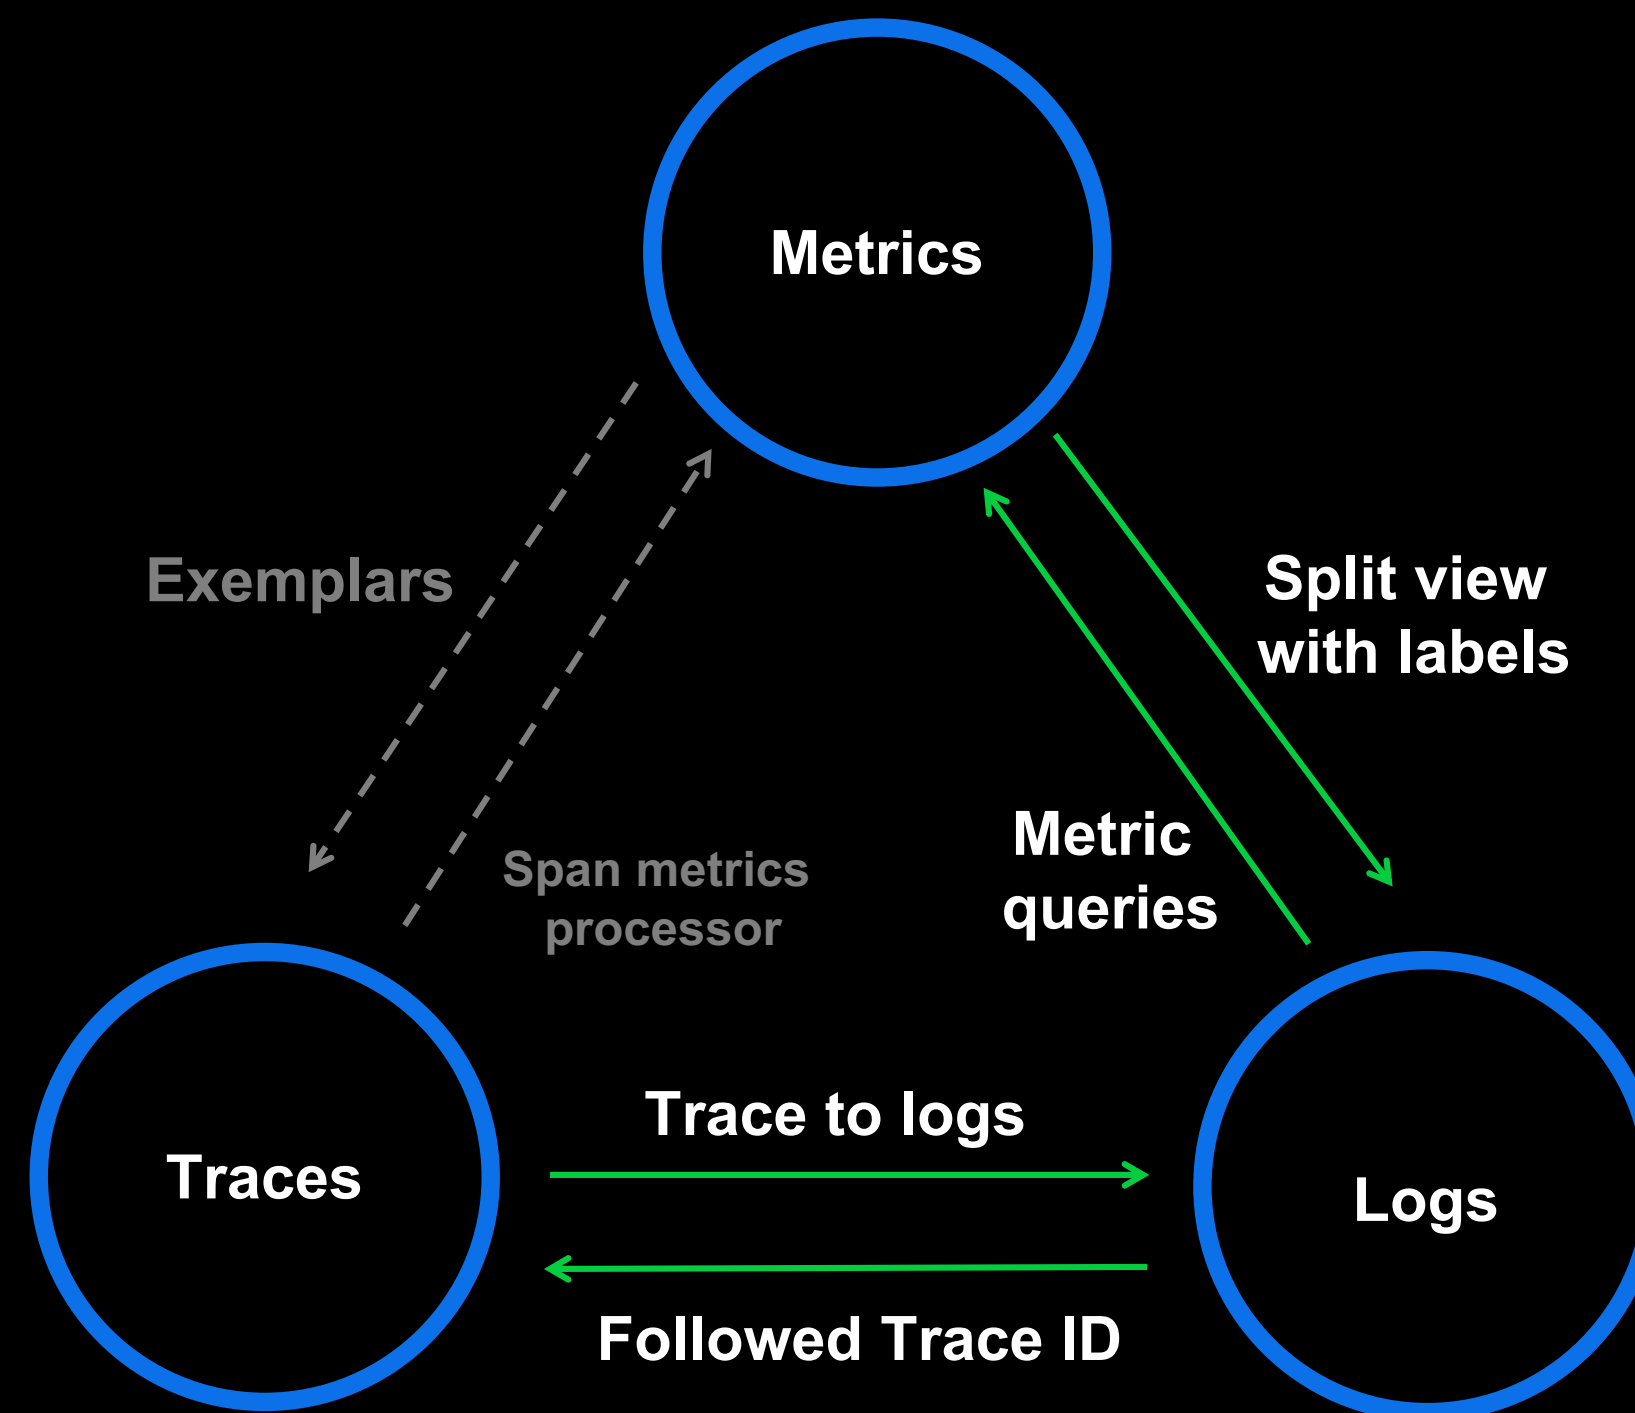
Inspect access logs

Traces

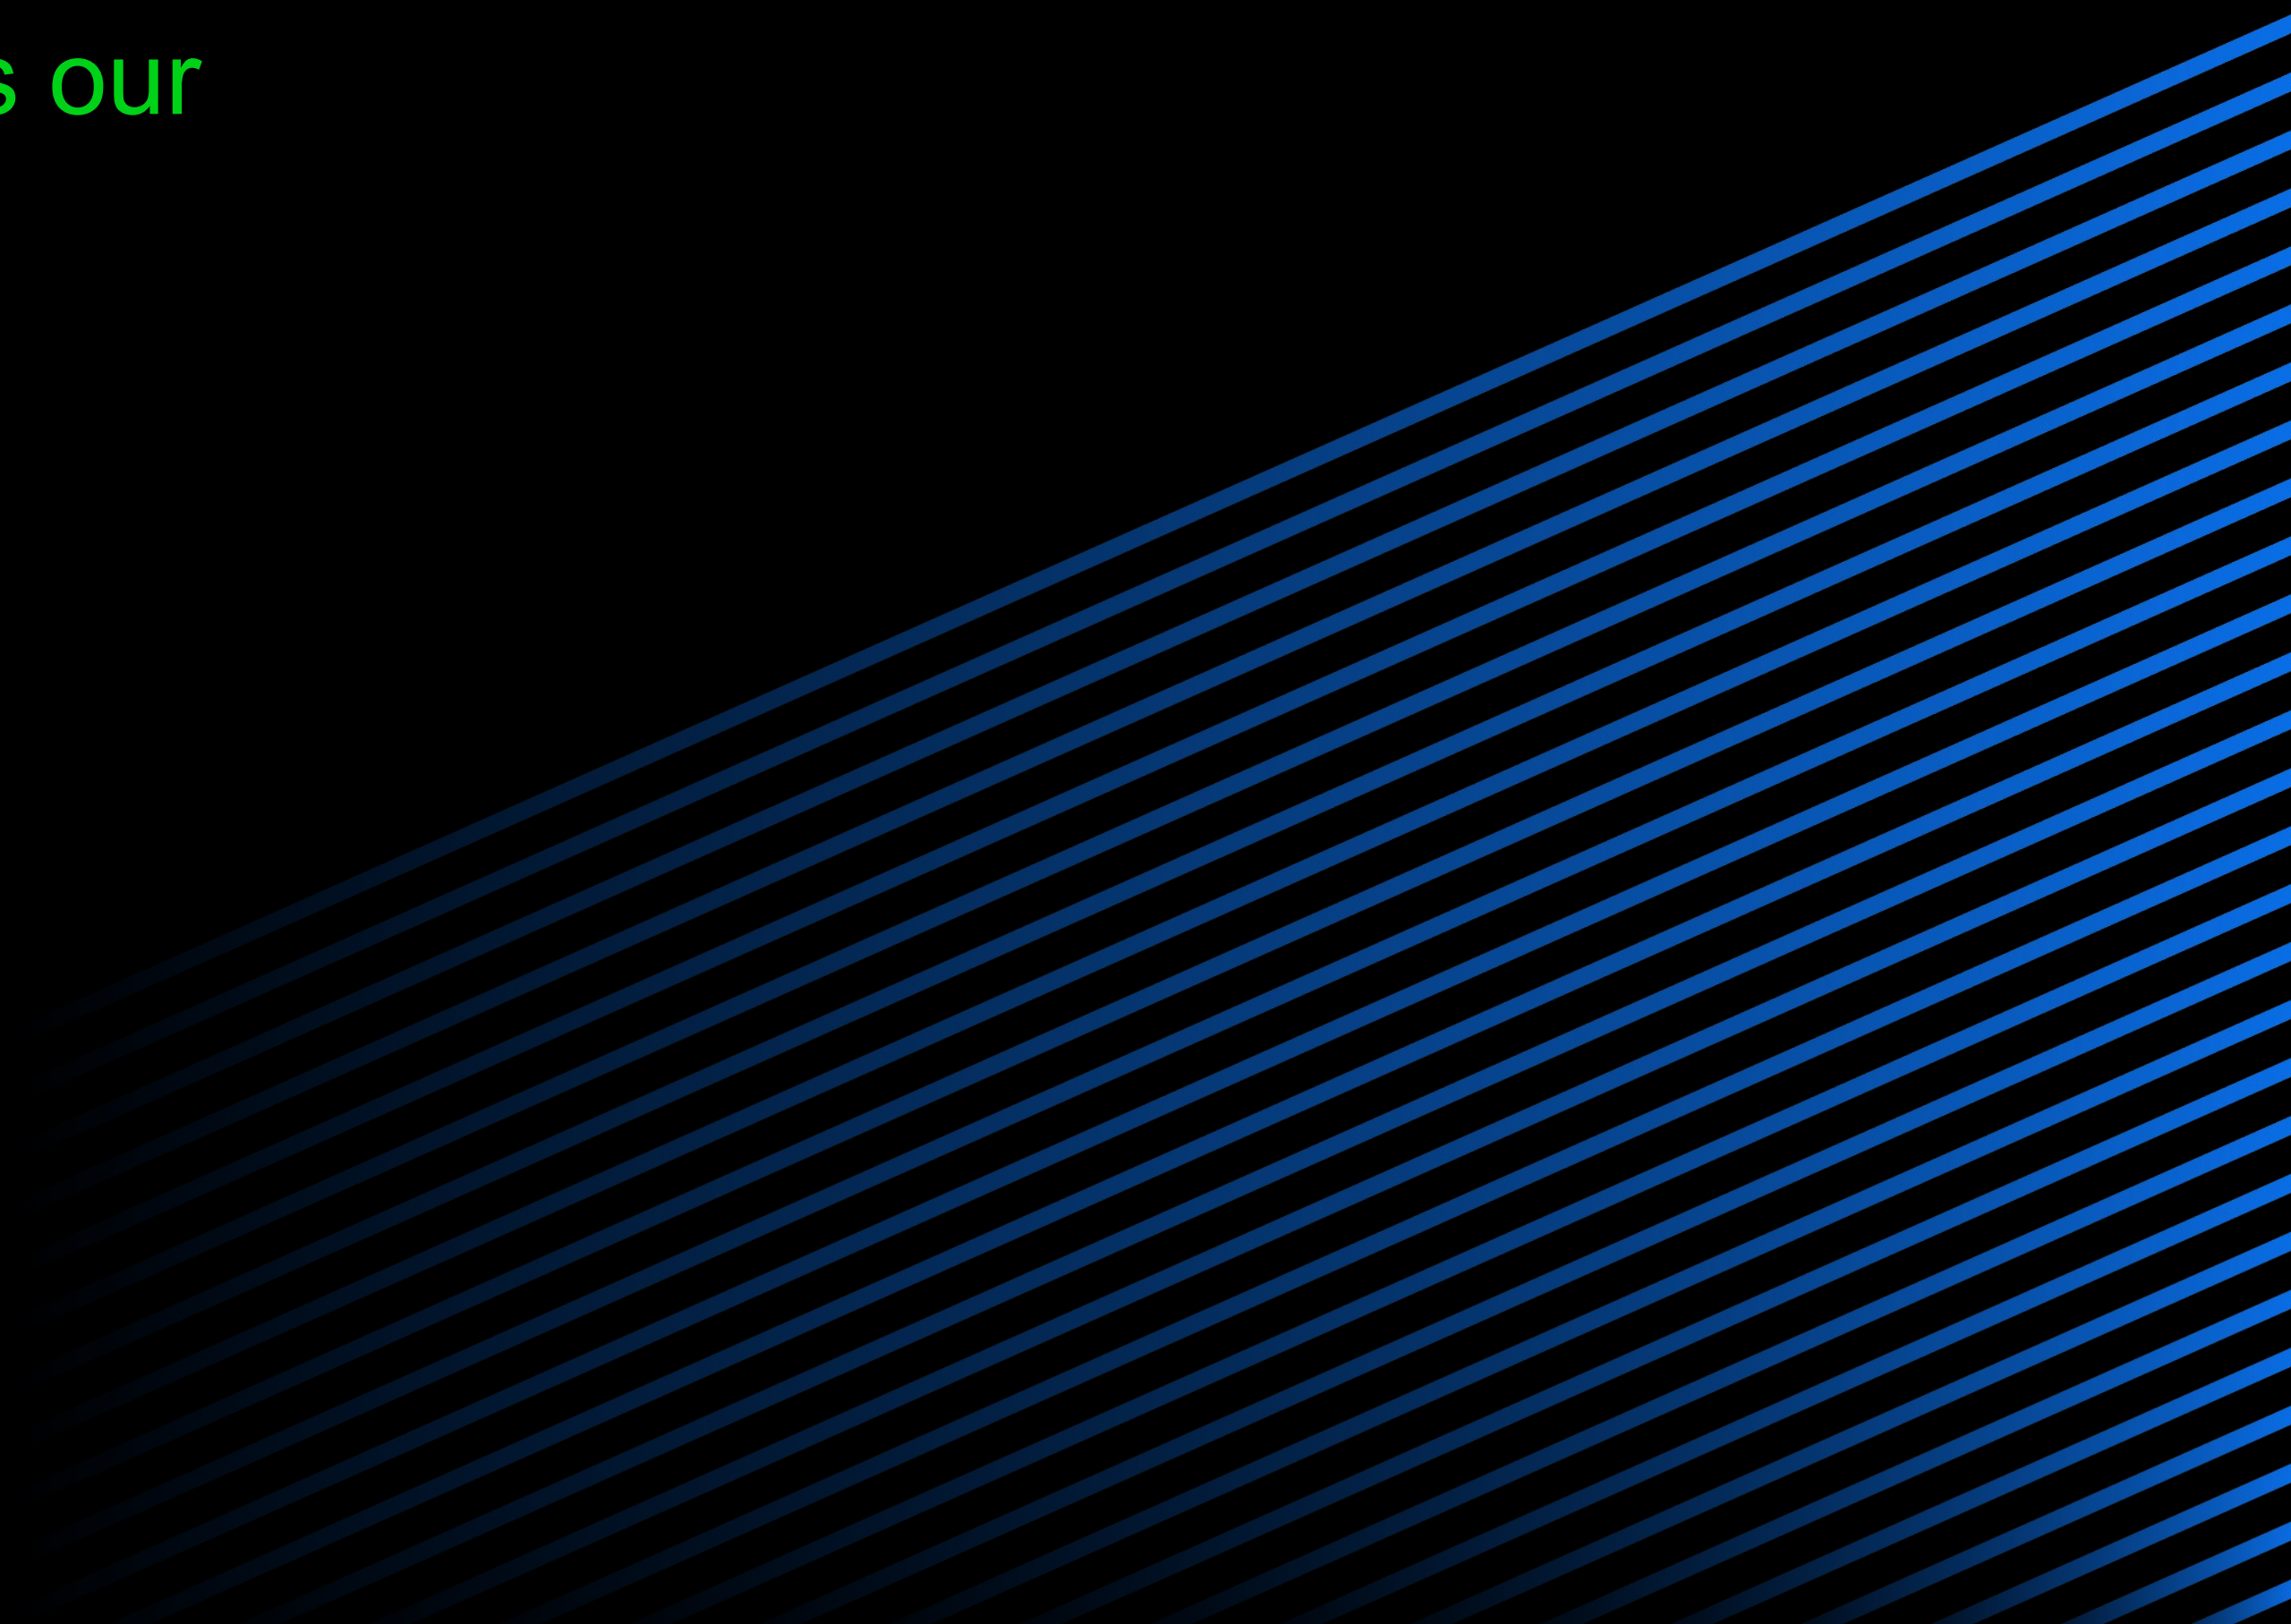
Open trace viewer

Logs

Jump to service logs of the trace



# Agenda

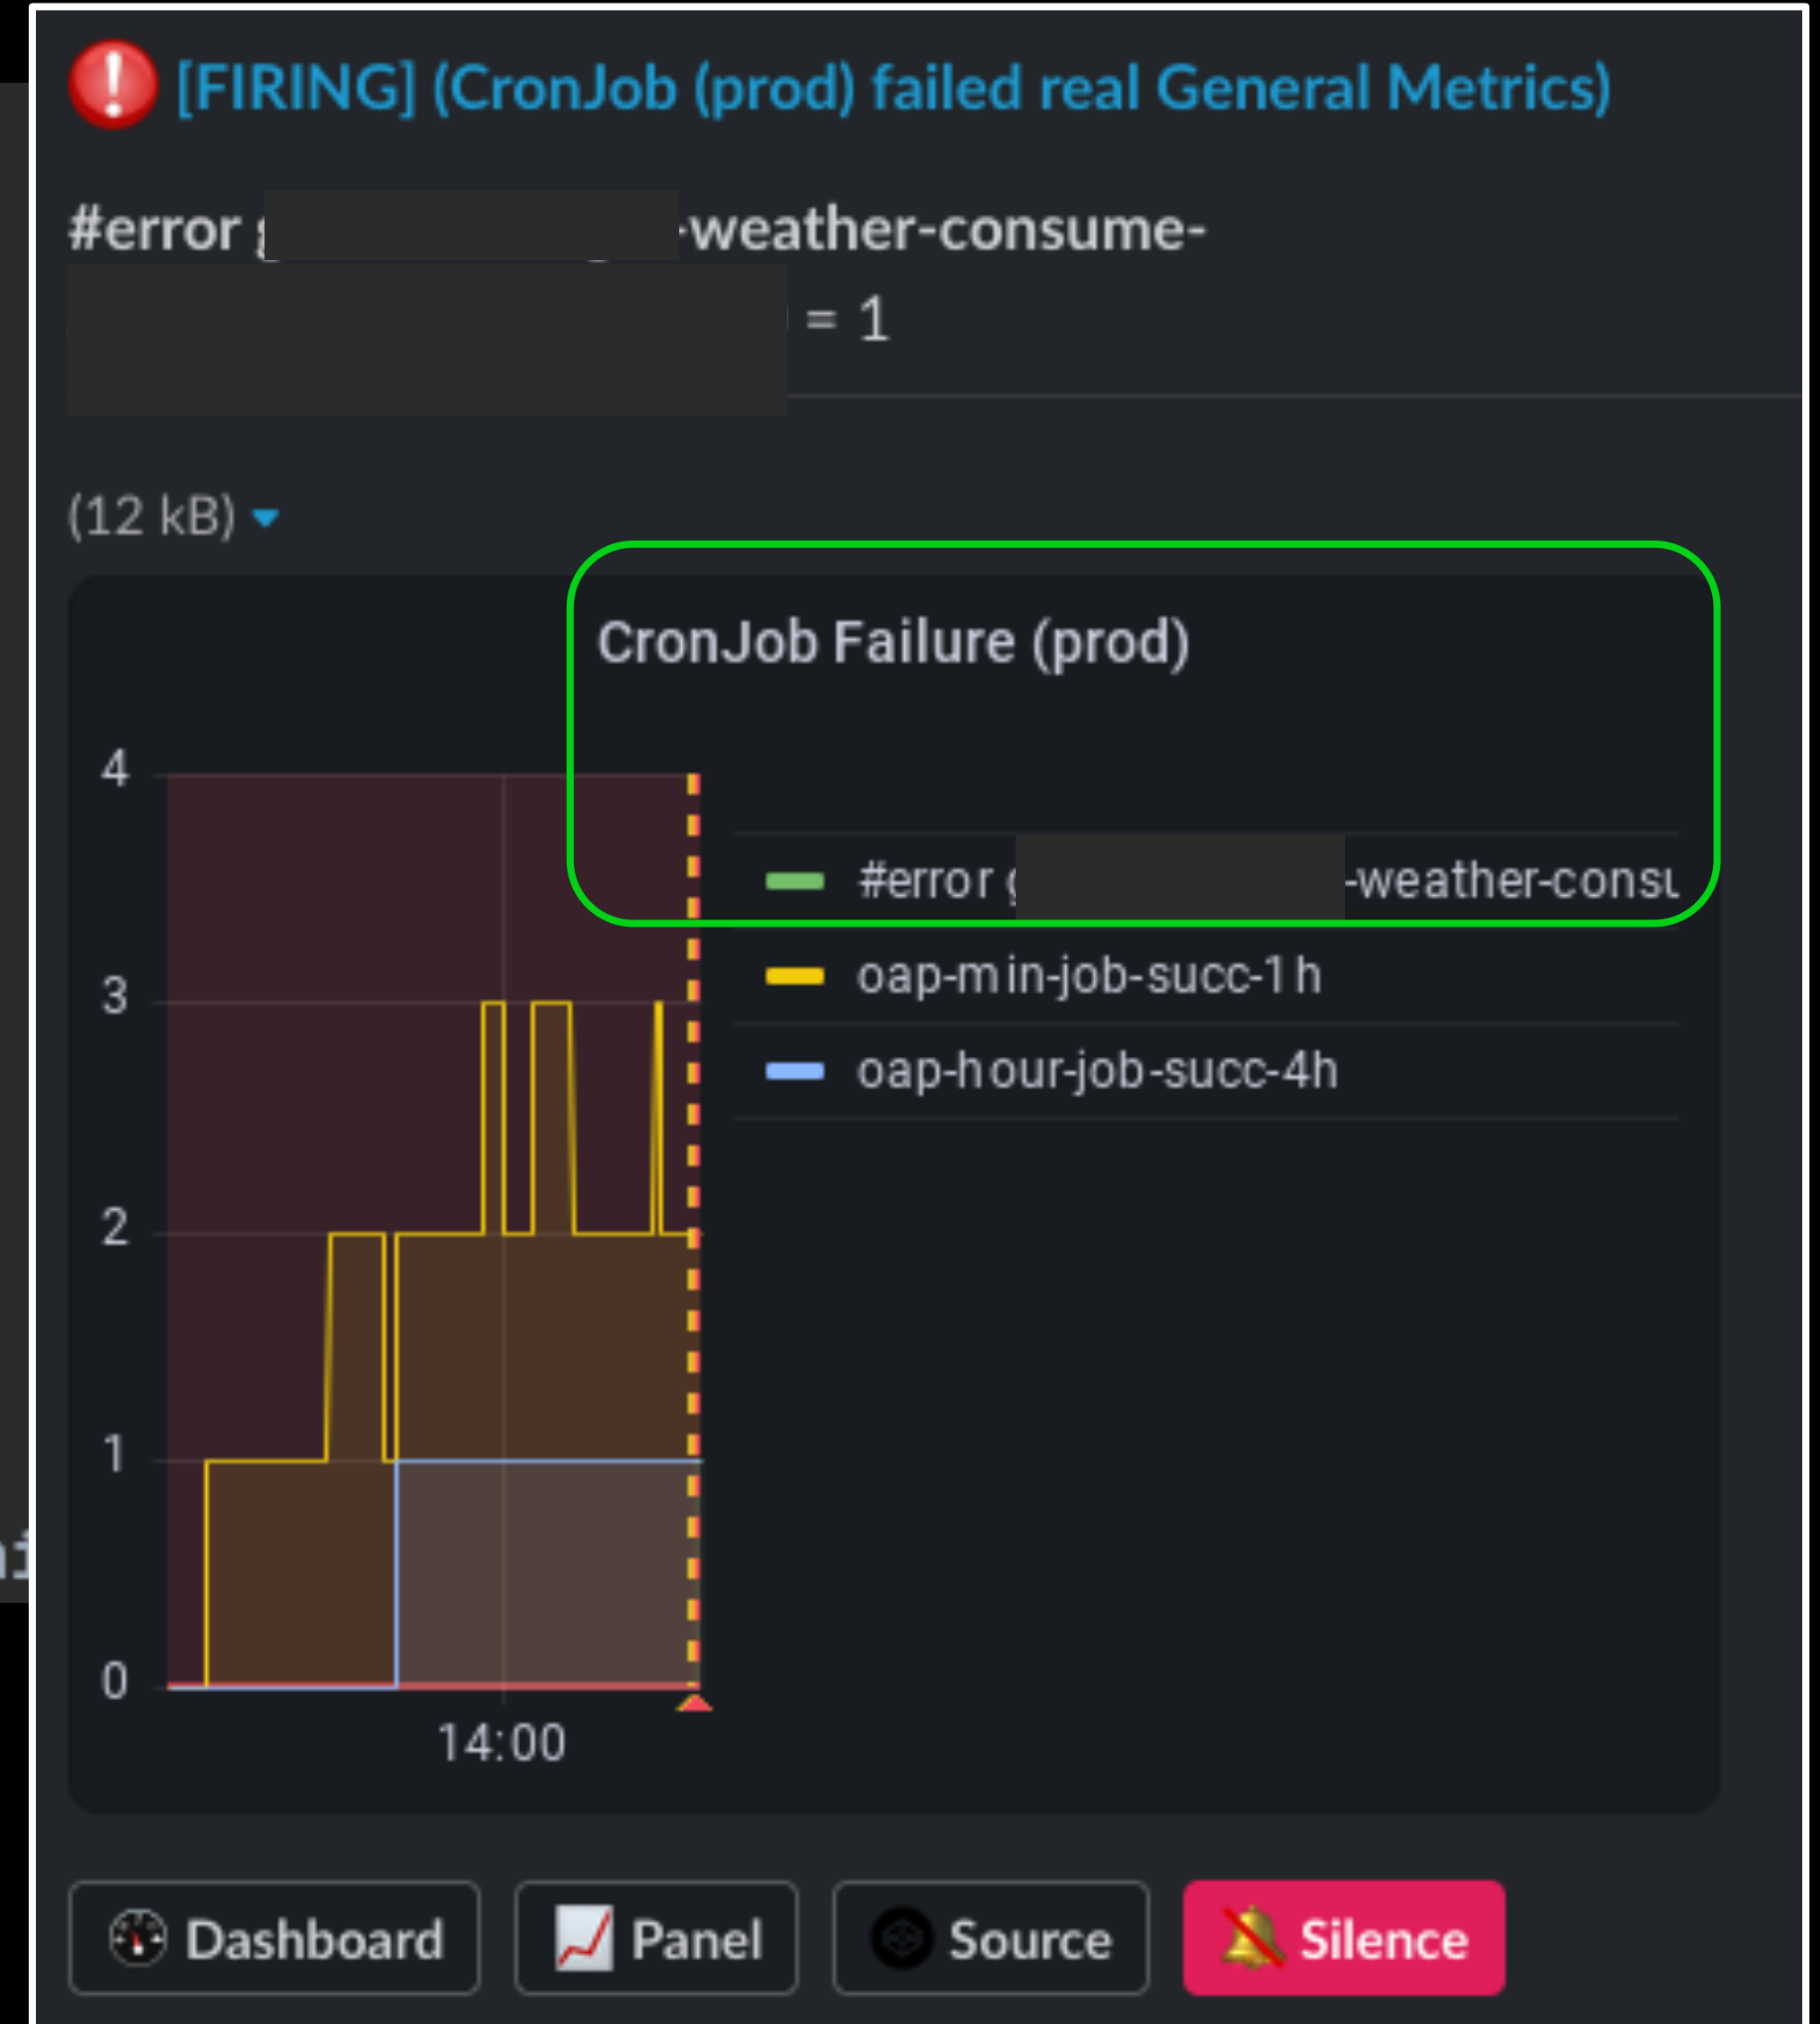
- LINE TODAY and its architecture
  - How miniservices and K8S changes our development and operation
    - Refactor to mini services
    - Refine CI/CD
    - Integrate with observability
    - Leverage K8S CronJob
  - Lessons learnt
- 

# Run periodic simple tasks (java) on Kubernetes

- **Requirements**
  - run at the specific time / interval
  - simple
  - **concurrency control**
  - **running history and logs**
  - **monitor**
  - **easy to run in local and test env**
- **Options**
  - **Spring @scheduled**
  - **Quartz**
  - **Spring Cloud Data Flow**
  - **AirFlow**
  - **K8S CronJob**

# Use K8S CronJob to run periodic tasks

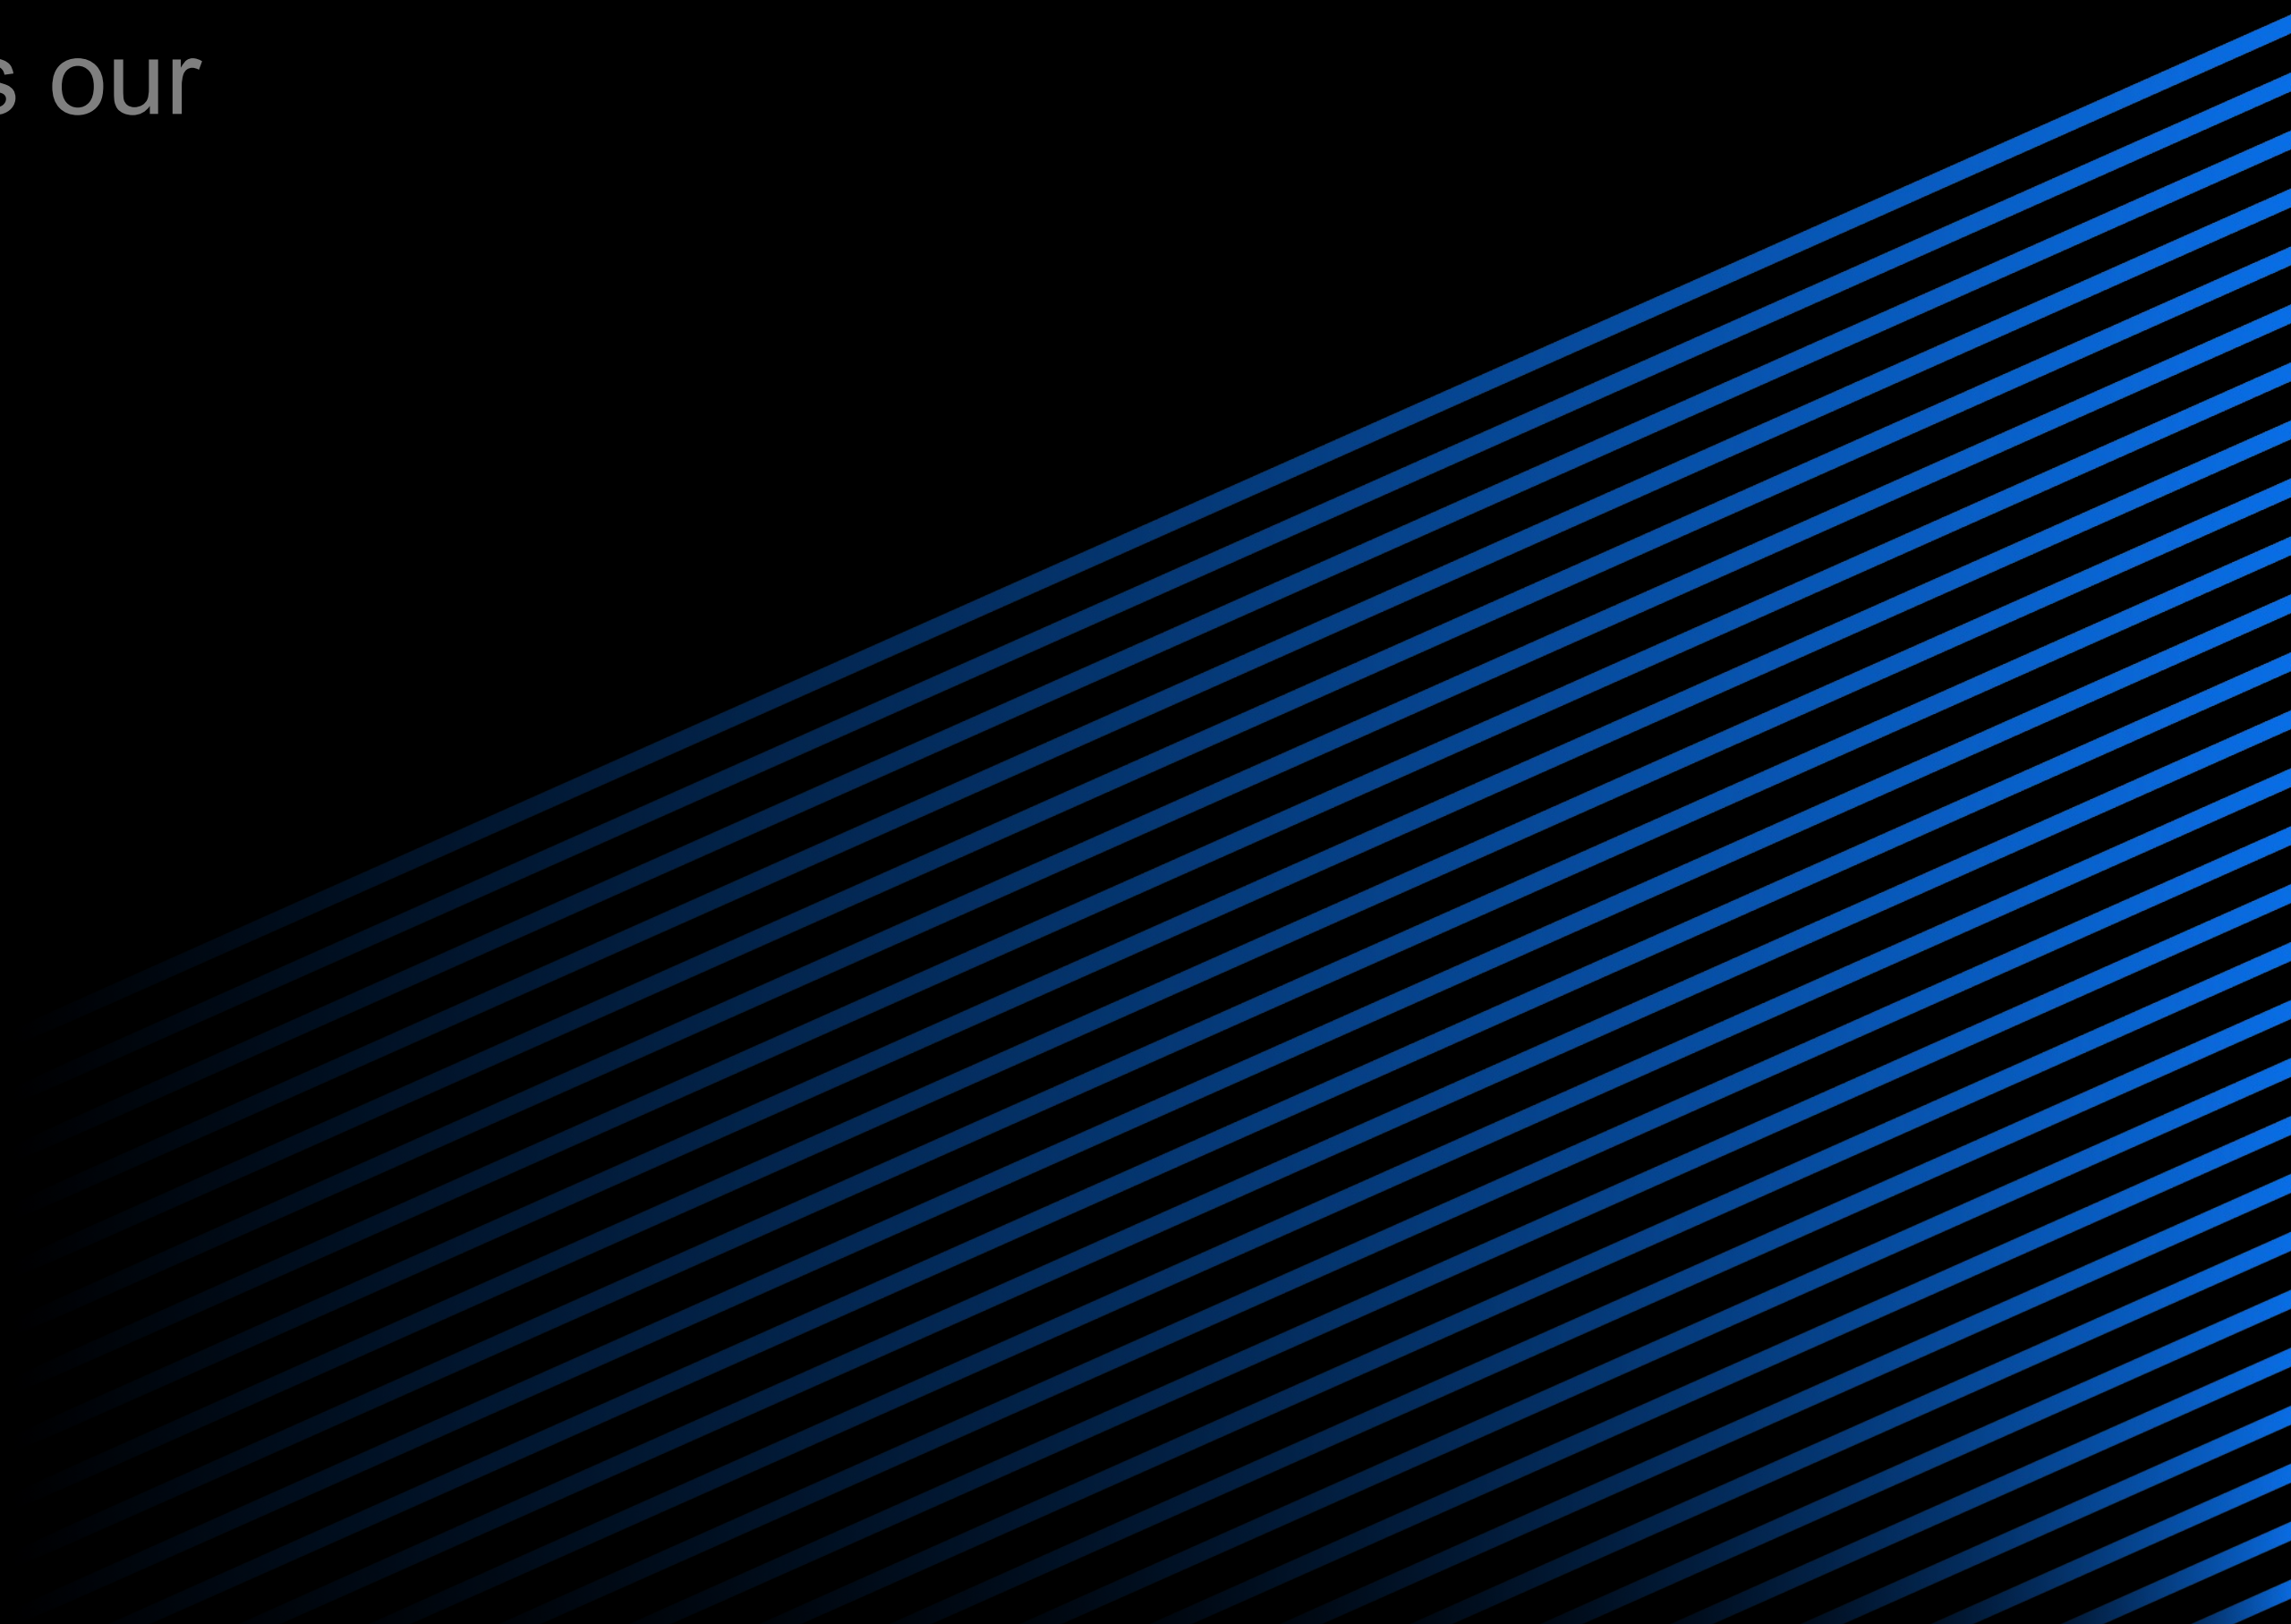
```
1  apiVersion: batch/v1beta1
2  kind: CronJob
3  metadata:
4    name: helloworld
5  spec:
6    schedule: "* * * * *"
7    successfulJobsHistoryLimit: 100
8    failedJobsHistoryLimit: 100
9    concurrencyPolicy: Forbid
10 suspend: false
11 jobTemplate:
12   spec:
13     template:
14       spec:
15         containers:
16         - name: helloworld
17           image: docker-registry.linecorp.com/line-ta
```



# K8S CronJob metrics

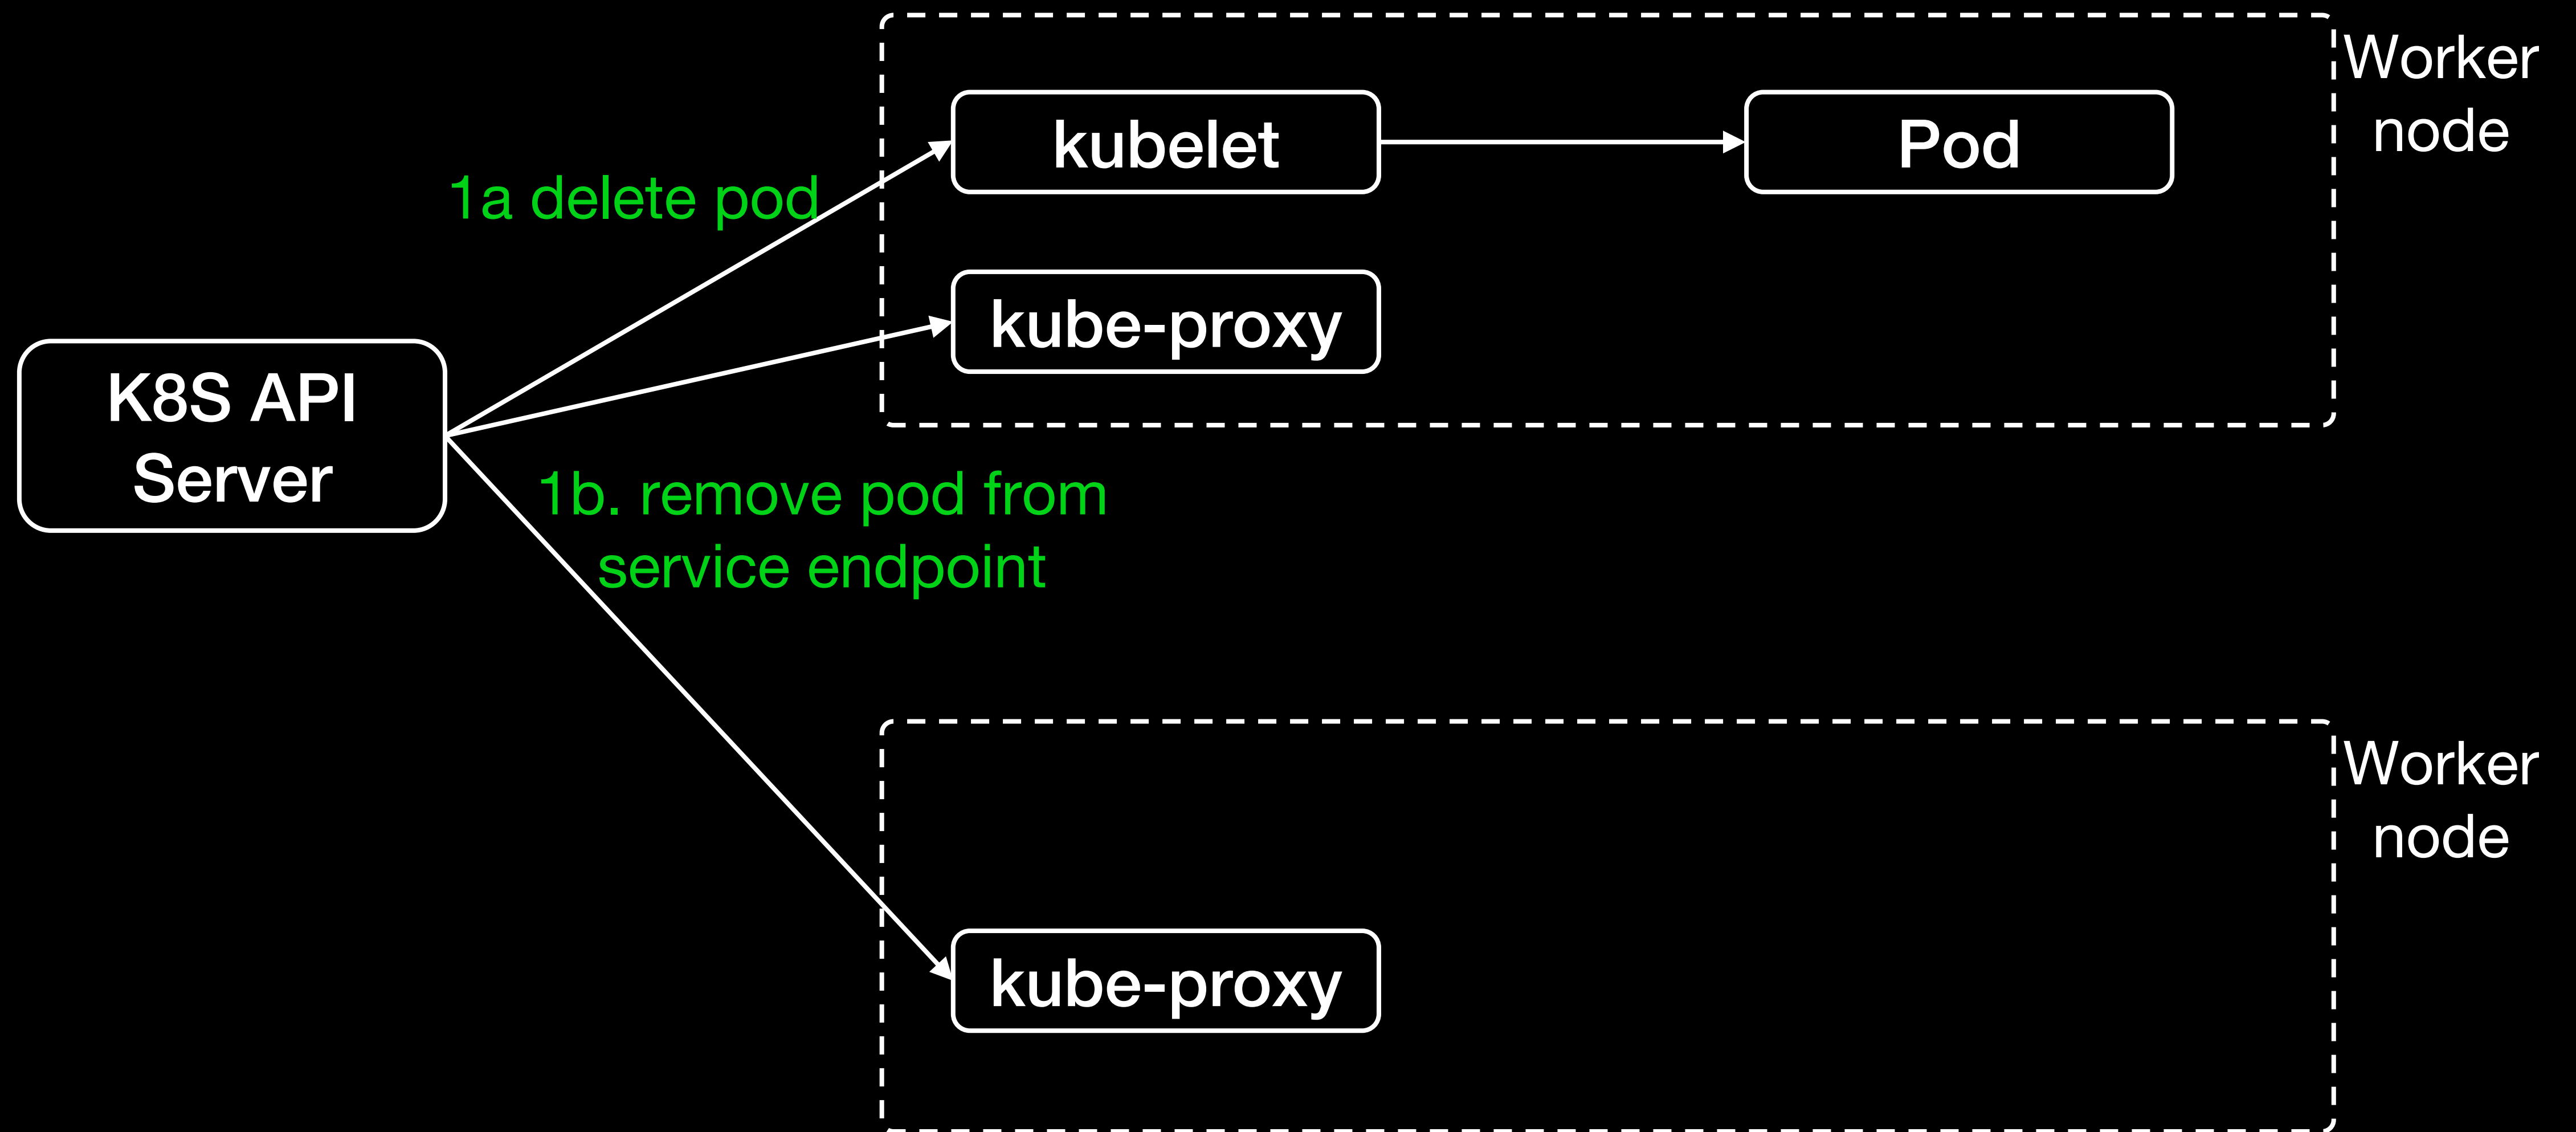
|                                   |  |                                 |  |  |
|-----------------------------------|--|---------------------------------|--|--|
| <p>Failed CronJob</p> <h1>OK</h1> | <p>Schedule</p> <p>40 17 * * *</p>       | <p>Failed Cronjob</p> <p>OK</p> | <p>Maximum Dur...</p> <p>1.45 min</p>    | <p>Failed Jobs Name</p> <p>No data</p> |
|                                   | <p>Next Schedule</p> <p>in 7 hours</p>   | <p>Active CronJob</p> <p>0</p>  | <p>Minimum Dura...</p> <p>1.45 min</p>   |  |
|                                   | <p>Last Schedule</p> <p>17 hours ago</p> | <p>Succeeded Cr...</p> <p>1</p> | <p>Average Durati...</p> <p>1.45 min</p> |  |

# Agenda

- LINE TODAY and its architecture
  - How miniservices and K8S changes our development and operation
  - **Lessons learnt**
- 
- A decorative graphic consisting of numerous parallel blue diagonal lines that sweep across the bottom right portion of the slide, creating a sense of motion and depth.



# Issue - intermittent errors during rolling update



# Solution - graceful shutdown

deployment manifest

```
lifecycle:  
  preStop:  
    exec:  
      command:  
      - sh  
      - -c  
      - sleep 10
```

spring boot application.yaml

```
server:  
  shutdown: graceful
```

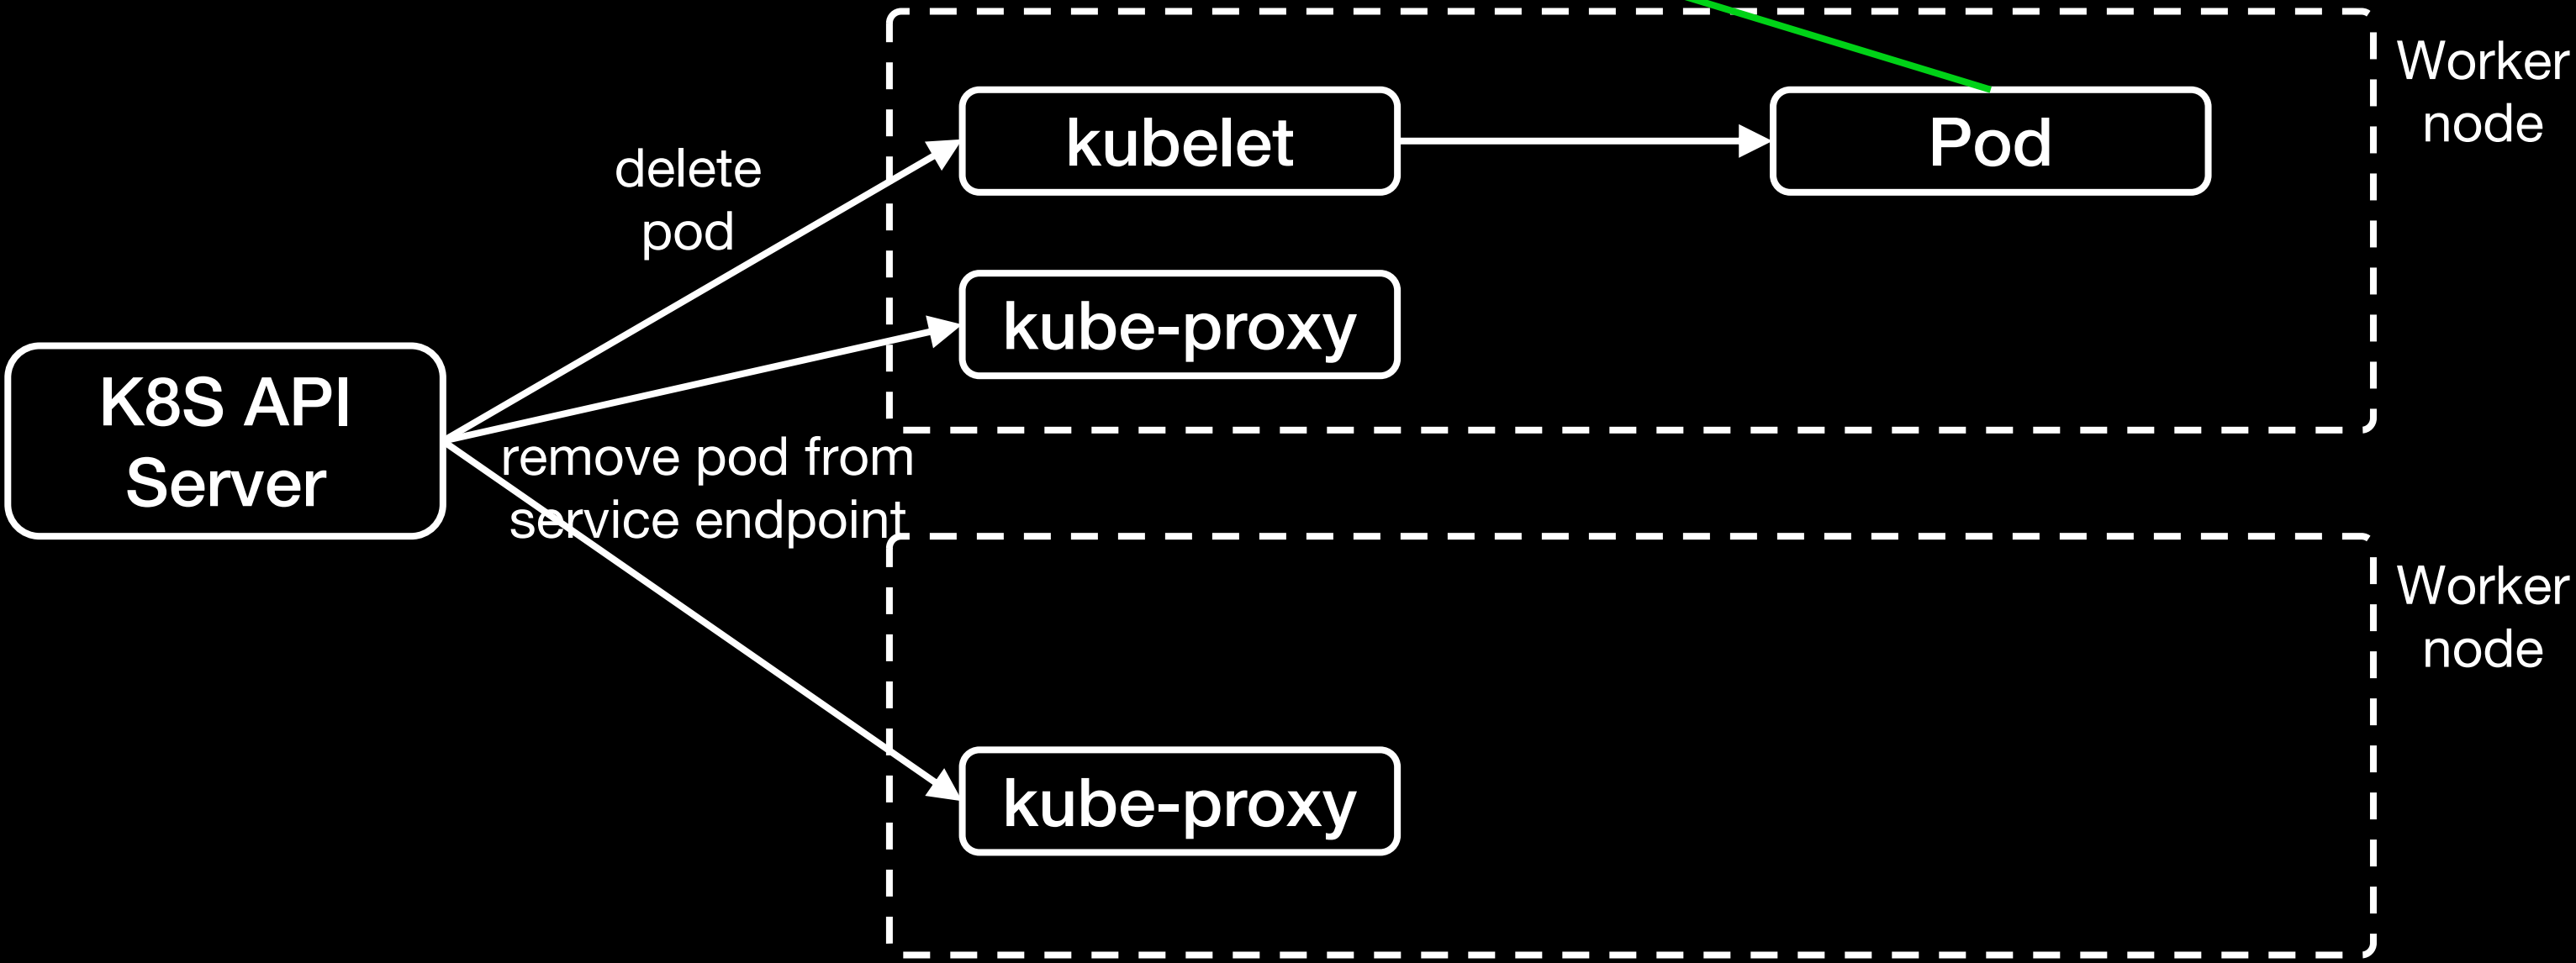
- Existing services allowed to complete
- No new requests permitted

Container killed (if running)

Main container process



Container shutdown

Pre-stop hook



# Issue - unpredictable request spikes

```
readinessProbe:  
  httpGet:  
    port: 8080  
    path: /health/readiness  
  periodSeconds: 10  
  timeoutSeconds: 1  
  failureThreshold: 3  
livenessProbe:  
  httpGet:  
    port: 8080  
    path: /health/liveness  
  periodSeconds: 20  
  timeoutSeconds: 1  
  failureThreshold: 3
```



- pod removed from endpoint at 30ish seconds
- fewer pods available to serve requests
- requests queue up
- pod restarted at 60ish seconds
- downward spiral

# Options to handle request spikes - it depends

- **Overprovision**
  - **\$\$\$\$**
- **Auto scaling**
  - **pod - 20+ seconds**
  - **node - ~5 minutes**
  - **serverless (lambda) - seconds**
- **Protection via ingress controller / api gateway**
  - **circuit breaker - 503**
  - **rate-limit - 429**
- **Improve design**




# Root cause: linux kernel memory leak

```
# cat /proc/meminfo
MemTotal:      8008704 kB
MemFree:       126904 kB
MemAvailable:  2402932 kB
Buffers:       60 kB
Cached:        12328 kB
SwapCached:    0 kB
Active:        900820 kB
Inactive:      9508 kB
Active(anon):  900432 kB
Inactive(anon): 336 kB
Active(file):  388 kB
Inactive(file): 9172 kB
Unevictable:   0 kB
Mlocked:       0 kB
SwapTotal:     0 kB
SwapFree:      0 kB
Dirty:         0 kB
Writeback:     0 kB
AnonPages:     898288 kB
Mapped:        3672 kB
Shmem:         2480 kB
Slab:          6750240 kB
SReclaimable:  2492932 kB
SUnreclaim:   4257308 kB
KernelStack:  16336 kB
PageTables:    11608 kB
```

reboot



 **jpmenil**

[Comment on #841 Docker fails to start containers with cgroup memory allocation error.](#)

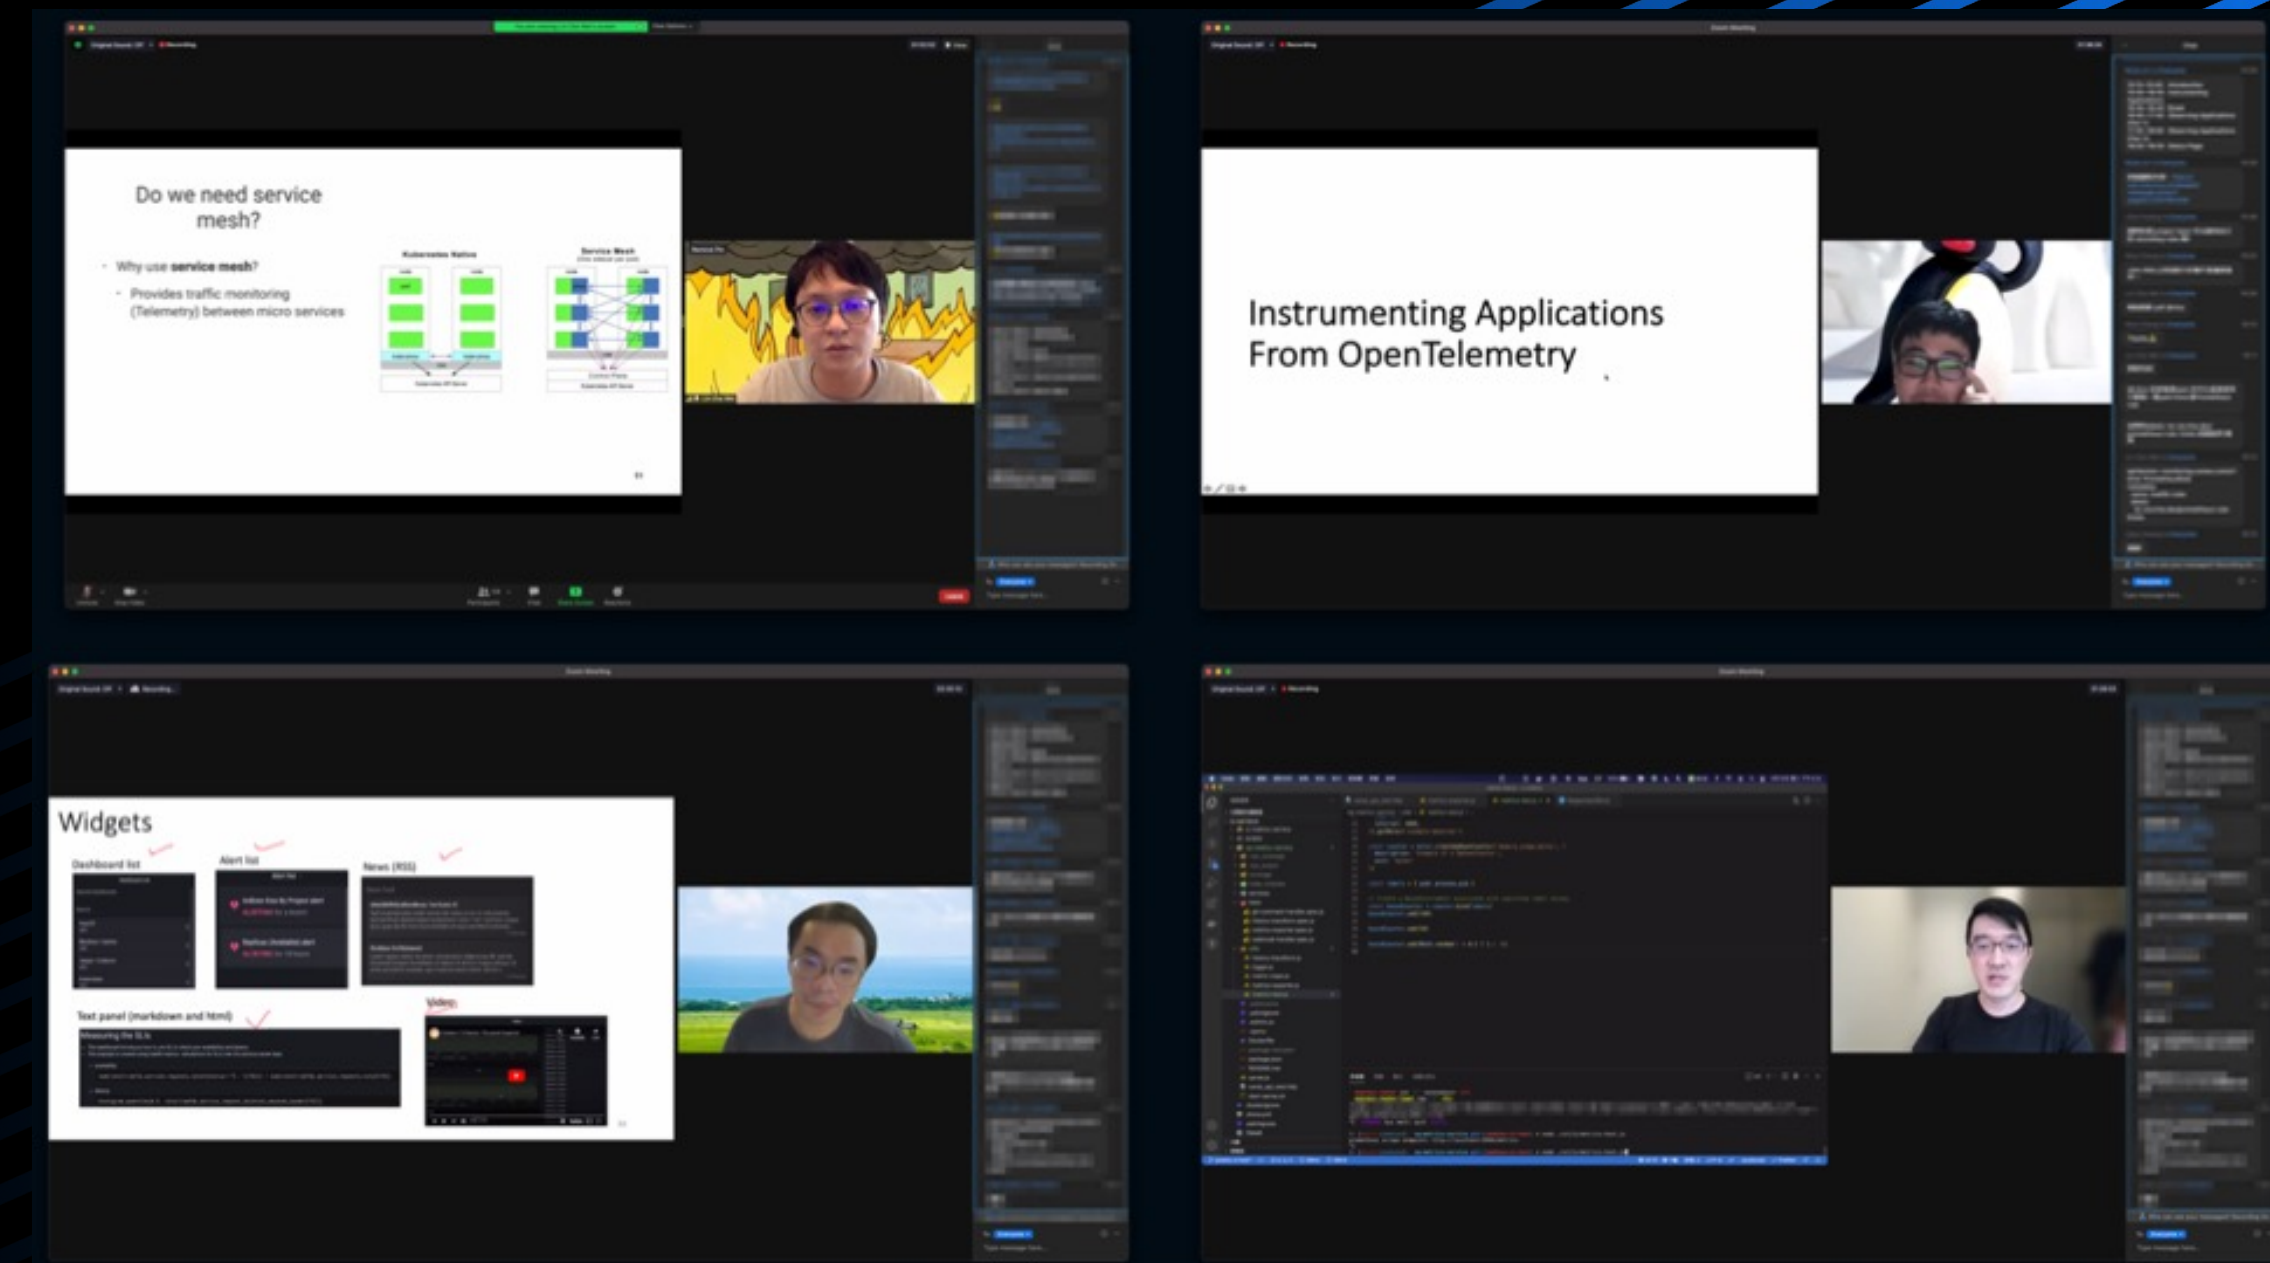
Must be fixed with kernel kernel-3.10.0-1075.el7

docker/for-linux | Dec 13th, 2019 | Added by GitHub

```
Inactive:      998908 kB
Active(anon):  1121036 kB
Inactive(anon): 340 kB
Active(file):  440280 kB
Inactive(file): 998568 kB
Unevictable:   0 kB
Mlocked:       0 kB
SwapTotal:     0 kB
SwapFree:      0 kB
Dirty:         8 kB
Writeback:     0 kB
AnonPages:     1119600 kB
Mapped:        218276 kB
Shmem:         1784 kB
Slab:          136192 kB
SReclaimable:  80400 kB
SUnreclaim:   55792 kB
KernelStack:  11520 kB
```

# Summary

- LINE TODAY and its architecture
- Mini services and K8S helps dev / ops efficiency for large systems
  - Refactor to mini services
  - Refine CI/CD
  - Integrate with observability
  - Leverage K8S CronJob
- Build in-depth DevOps and K8S skills



Thank you

