

CNS411

# Optimizing Kafka workloads with AWS Lambda



**Alexander Vladimirov**  
Senior Solutions Architect, Serverless  
AWS

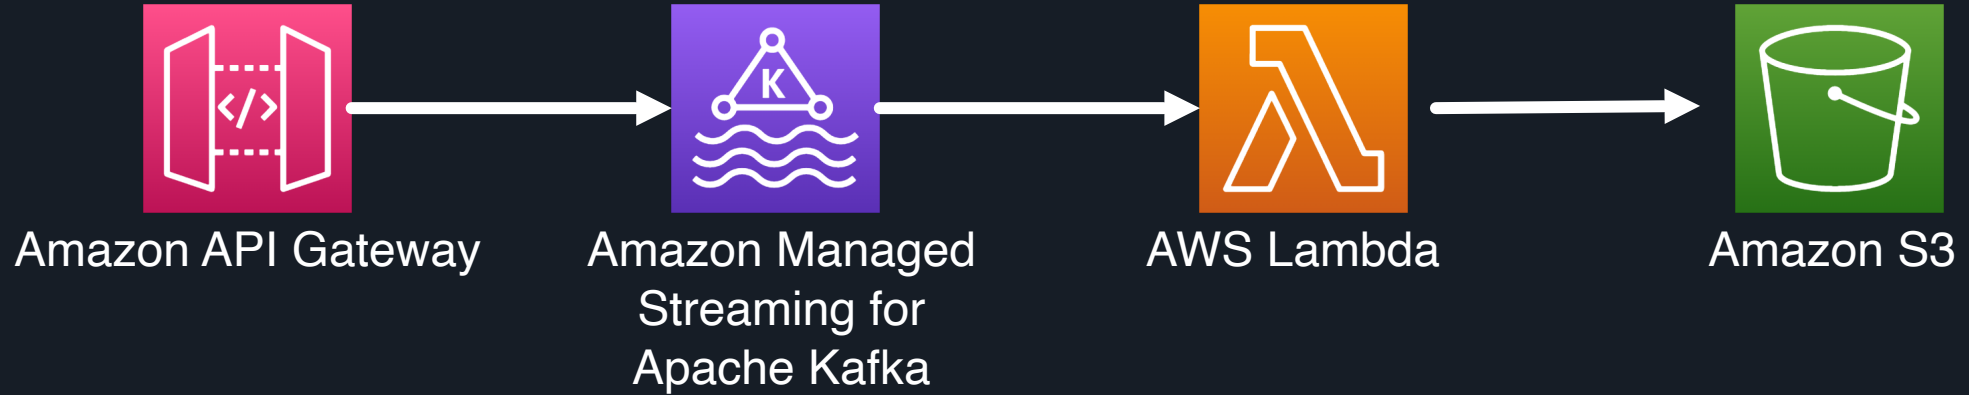


**Anton Aleksandrov**  
Principal Solutions Architect, Serverless  
AWS

# The business scenario – booking flights



# What does your architecture look like?



Throughput

End to end latency

Additional processing

Traffic patterns

# Event Source Mapping for Kafka

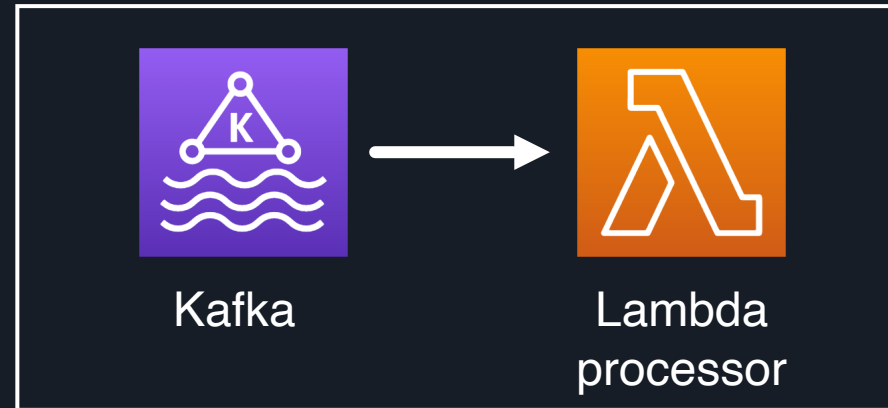
**Handles polling and automatically invokes Lambda functions**

**Supports batching, error handling and retries**


**Automatically scales and manages concurrency**





# Usecase 1 – baseline (out of the box defaults)



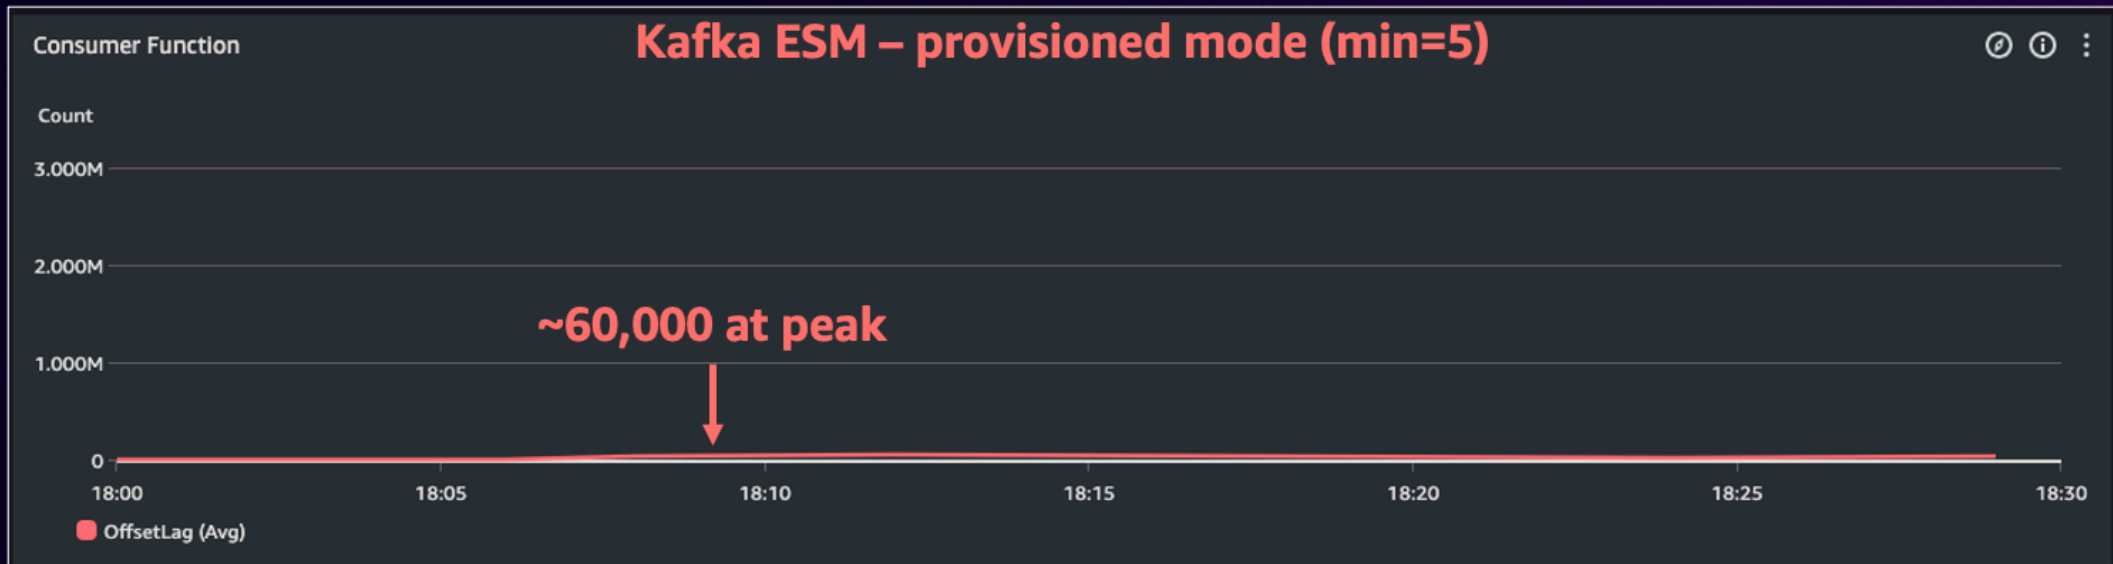

# Usecase 2 – Latency sensitive data streaming




# Usecase 2 – Latency sensitive data streaming



# Usecase 2 – Latency sensitive data streaming



# Usecase 3 – Even higher throughput with SQS




**Demo time!**


# Best practices and considerations




**Understand scaling behavior and quotas**



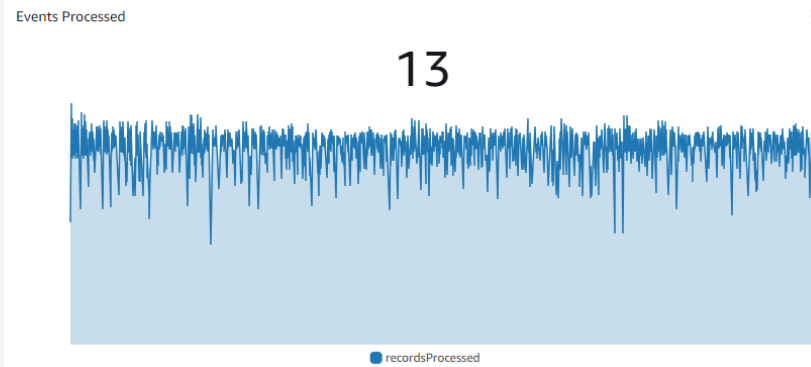
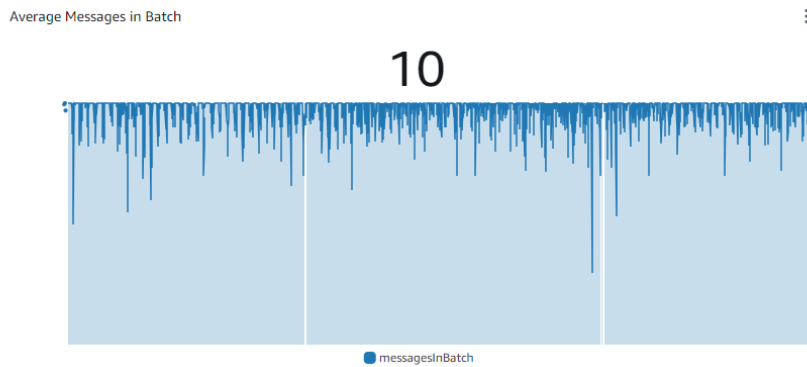


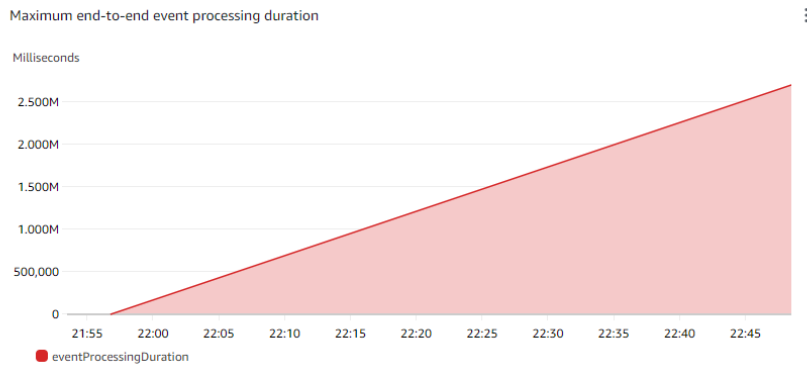
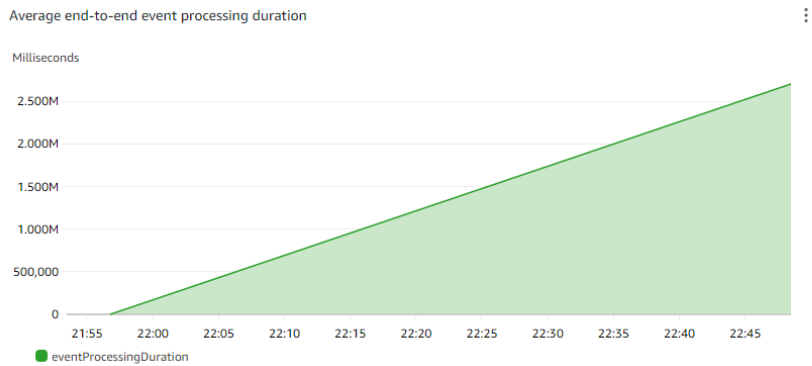
**Use batching, filtering, parallelization**




**Use provisioned mode when need low latency guarantees**




**Optimize with Lambda and SQS**




Average end-to-end event processing duration



Maximum end-to-end event processing duration




Minimum end-to-end event processing duration




Events Published

300




Average Messages in Batch

17.5



Events Processed

700




Events Published

700




Average Messages in Batch

1



Events Processed

700



# Conclusion

# 1

Know the out-of-the box numbers. Default processing works for majority of common use-cases.

# 2

Batching and parallelization are your best friends. Use them wherever possible.

# 3

SQS can help to further increase throughput for high-volume or spiky workloads.

# Check out these other sessions

## **API305: Building EDAs with Kafka and EventBridge**

**Monday (Dec 1) @ 1:30pm – Mandalay Bay – South Seas H – Chalk talk**

**Tuesday (Dec 2) @ 5:30pm – MGM – Boulevard 167 – Chalk talk**

## **CNS333: Serverless data processing with Lambda and Kafka**

**Wednesday (Dec 3) @ 5:30pm – MGM 104 – Builders Session**

## **CNS423: From trigger to execution: The journey of events in Lambda**

**Wednesday (Dec 3) @ 8:30am – Mandalay Bay – South Seas F – Breakout**



# Serverless & Appt Resources




<https://s12d.com/RIV2025>

- Continue your AWS serverless learning
- Powertools for AWS Lambda
- Serverless Land patterns

## Serverless **Patterns** Collection

Build integrations using infrastructure as code with **serverless patterns**





# Thank you!

**Alexander Vladimirov**

 avladi

**Anton Aleksandrov**

 antonal80

 [aal80.github.io/whoami](https://github.com/aal80/whoami)



This session's content  
and much more

Please complete the session  
survey in the mobile app