WEBINAR Building Streaming and Fast Data Applications with Spark, Mesos, Akka, Cassandra and Kafka

Sean Glover (@seg1o), Senior Consultant at Lightbend



A Bit of History: IT Services vs. Big Data



The Recent Past...

Services **Big Data**

Some Overlap: Concerns, Architecture



The Future?

Microservices & *Fast* Data

Much More Overlap

Drilling Down: Monoliths to Microservices



lightbend.com/reactive-microservices-arc

O'REILLY'

Reactive Microservices Architecture

Design Principles for Distributed Systems



Jonas Bonér



Nonoliths

 Tangled responsibilities, lead to infrequent, "big-bang" deployments • App lifetimes months to forever!

Lord	Balancer
------	----------





O'REILLY'

Reactive Microservices Architecture

Design Principles for Distributed Systems



Jonas Bonér

Microservices

 Each does one thing, so must be message driven & asynchronous • Updates easier, deployments frequent • App lifetimes: minutes! to forever



Microservices

But be careful: Message overhead is much better than function calls!



Reactive Systems



resilient

Pivotal

The Reactive Platform

Drilling Down: Big Data

The Emergence of Fast Data (Time Is Money)

Respond to change

The New York Times

Fast Data and Microservices; Are they Converging?

Synergies

• Each [stream app or µservice]: does one responsibility ingests unending [data or messages]

Each [stream app or µservice] must: operate asynchronously offer never-ending service

Synergies

Thesis

 These architectures are converging: 1.Similar design problems 2. Data becomes dominant problem

Value Three Ways

1. An accelerated on ramp for building streaming data systems, data applications, and other microservices.

Value Three Ways

2. Best practices guidance for solving specific design problems: - Sample apps - Documentation - Enablement services

Value Three Ways

3. Machine learning-based monitoring and management: -Keep your systems resilient, scalable, and responsive with minimal user intervention.

Service 1

Service 2

Service 3

Consumers

Service 1

Service 2

Service 3

Consumers

Analysis Machine Learning

Streaming Tradeoffs (1/3)

 Low latency? How low? High volume? How high?

Streaming Tradeoffs (2/3)

 Which kinds of data processing & analytics are required? How will this processing be done? Individual processing of events? Bulk processing of records?

Streaming Tradeoffs (3/3)

 Which tools and data sources/ sinks must interoperate with your streaming tool?

 Low latency Low volume Complex flows Complex Event Processing

• Med. latency High volume • Data flows, SQL • En masse processing

 Low latency High volume • Data flows, correctness • En masse processing

Lightbend.ai Cluster Analysis Machine Learning

Cloud Hosted

Kafka & Spark Metrics

Correlated Troubleshooting

Why Only Hosted Services?

• Only metrics are uploaded, not sensitive domain data. One less service for you to manage.

Why Only Hosted Services?

• We can rapidly evolve this service. without impacting your environment. You benefit from aggregated knowledge from all FDP clusters.

Upgrade your grey matter!

Get the free O'Reilly book by Dr. Dean Wampler, VP of Fast Data Engineering at Lightbend

bit.ly/lightbend-fast-data

