THE DEVELOPER'S CONFERENCE

Using a Serverless Architecture to Deploy (and consume) Machine Learning models

Rafael Zotto
Senior Software Architect, HP Inc.

Short Bio

Rafael Zotto

Holds a master degree in Computer Science focused in high performance computing. Specialized in parallel and distributed computing with special interest in mobile and web technologies. Works for HP Inc. for the past decade acting as senior software architect for print firmware and wearable technologies. Recently joined the Data Science research team in Porto Alegre, Brazil.

Agenda

THE DEVELOPER'S CONFERENCE

- Problem Statement
- Serverless Overview
- Deploy and Consume ML: A practical Use Case

Problem Statement

THE DEVELOPER'S CONFERENCE

- The model is trained;
- The results are acceptable;

- How to share it '*with the world*'?

  - Talk goal: to share a ***Serverless Approach***

Process Perspective

THE DEVELOPER'S CONFERENCE

- Development and Deployment of ML systems **should not be different** from traditional software solutions.

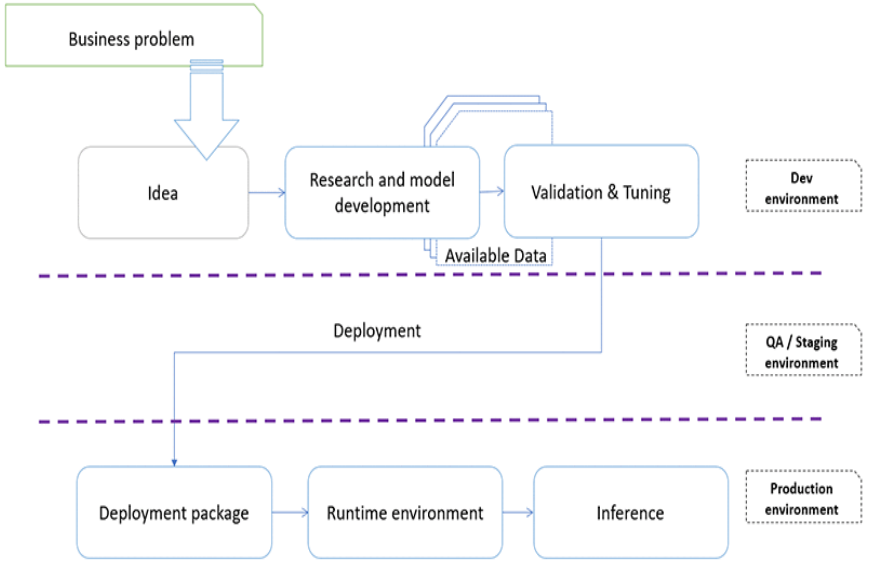

* Image from "*How to Deploy Deep Learning Models with AWS Lambda and Tensorflow*" article. Amazon Web Services Blogs.

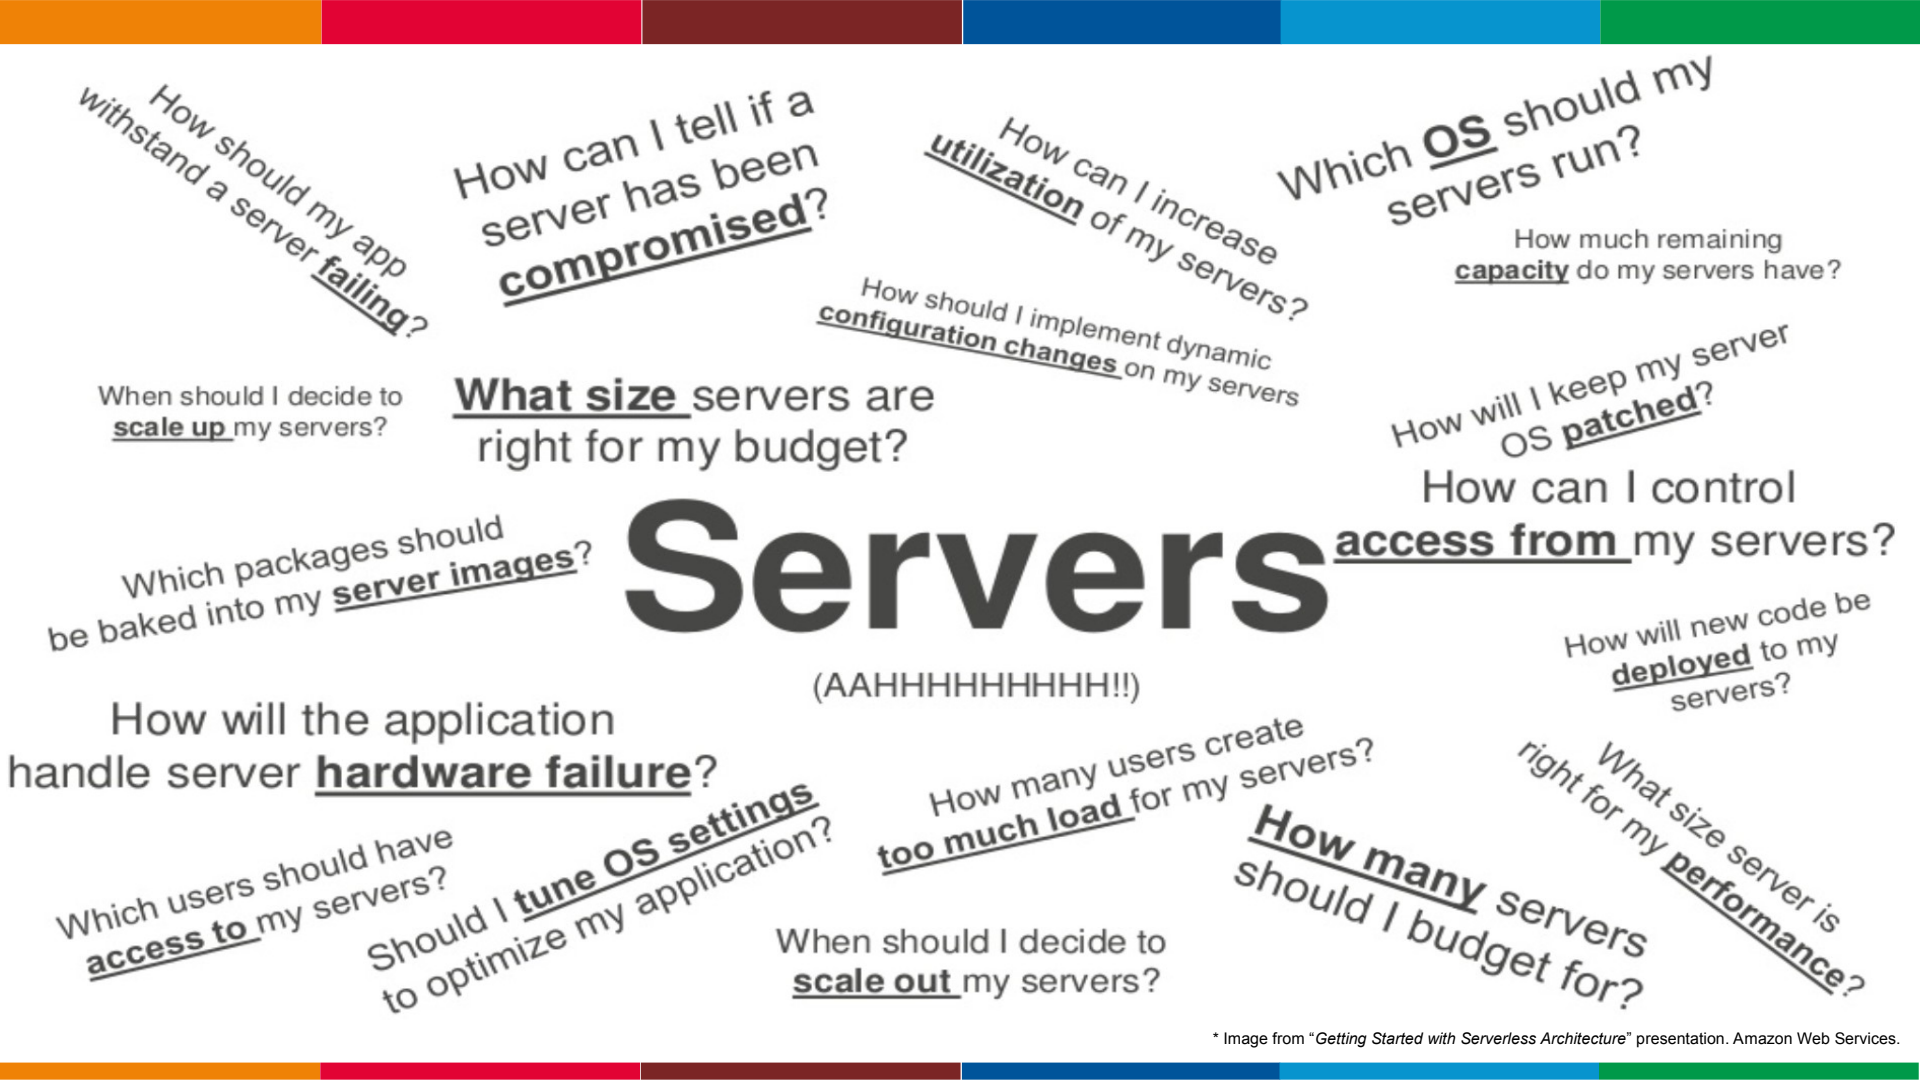

* Image from "*Getting Started with Serverless Architecture*" presentation. Amazon Web Services.

Serverless Definition

THE DEVELOPER'S CONFERENCE

- Platform to develop, run and manage applications without the complexity of building and maintaining infrastructure.

- No free lunch!
  - You will pay for it.
  - Sub-second billing

Architect to be Serverless

THE DEVELOPER'S CONFERENCE

- Fully Managed
  - No provisioning, zero administration, high-available
- Developer Productivity
  - Focus on what matters, innovate quickly
- Continuous Scaling
  - Up and Down automatically

Simple Use Case

THE DEVELOPER'S CONFERENCE

- **Model** was previously trained;
- **Deploy** it to *a* cloud environment;
- **Execute** real-time predictions;

Simple Use Case

THE DEVELOPER'S CONFERENCE

- **Model** – IRIS Data set
- **Deploy** – AWS
- **Execute** – AWS IoT Core Stack;

Model

THE DEVELOPER'S CONFERENCE

- The **Iris flower data** set is a **multivariate data set**
- Introduced by the British statistician and biologist Ronald Fischer (1936)

Iris Setosa

Iris Versicolor

Iris Virginica

Deployment

THE DEVELOPER'S CONFERENCE

- AWS Lambda functions for predictions
- Model saved in a S3 bucket

- The **Serverless Framework** might be your friend here.

Real-Time Prediction

THE DEVELOPER'S CONFERENCE

- IoT Core stack (just because we want it 'real-time')
- MQTT Communication (Lambda ↔ Client)

ENOUGH TALKING . . .
SHOW ME THE CODE.

Tips and Issues

THE DEVELOPER'S CONFERENCE

- Take advantage of AWS Lambda **container reuse**
- Keep you function **warm**!

THE DEVELOPER'S CONFERENCE