



Leveraging Serverless to Launch Applications Faster

Presenter: Chinmay A. Anaokar

This workshop offers an introductory demonstration on leveraging Serverless technology in AWS to quickly launch applications at a reduced cost. Many people struggle to make their applications live because they don't know—or don't want to deal with—setting up the infrastructure. This workshop aims to assist those individuals by providing a basic overview of Serverless technology and how it eliminates the need for server management. The intended audience is anyone interested in Serverless technology or software development. Attendees are advised to bring laptops for maximum benefit. The workshop covers a small project wherein the attendees will be able to launch their own application using Serverless. After attending this workshop, participants will be able to host applications with Serverless technology in the cloud, without launching complex infrastructure or incurring significant costs.

Requirements:

- Attendees are advised to bring their laptops.

Workshop Capacity: 20 participants



The Power of Profiling: Optimizing ML Workloads to Reduce Energy Usage

Presenter: Michael Shin

The resource need for machine learning (ML) product development is rising exponentially as state-of-the-art models are growing exponentially. These models require significant hardware resources to train. Optimizing hardware usage can help reduce costs, shorten development time, and reduce energy usage, leading to lower environmental impacts. To address these needs, my colleagues and I at CentML have created an open-source ML profiling tool called DeepView. In this workshop, I will showcase our tool and promote its adoption for promoting environmental sustainability.

Participants will install DeepView on their personal computers and will follow along with a demonstration that will explore DeepView's features and capabilities. If they cannot install DeepView on their personal computers, we will provide virtual machines that participants can access remotely. This workshop will show DeepView's accuracy and highlight how it can optimize hardware usage to reduce energy usage. Additionally, we will discuss how we can use the tool to select cloud providers that minimize their carbon impact.

The goal of showcasing DeepView is to increase its adoption and promote environmental sustainability. By measuring hardware usage during ML model training, DeepView can help optimize hardware utilization for ML workloads and reduce energy usage, leading to a reduction in carbon emission. Through this demonstration, we hope to create



awareness of the importance of energy conservation when it comes to ML development and the role of technology in promoting environmental sustainability. DeepView can be valuable to developers working with ML models, leading to sustainable and energy-efficient computing practices.

Requirements:

- Attendees **must** bring their laptops.

Workshop Capacity: 30 participants