

Towards Policy-aware Edge Computing Architectures

Pratik Baniya, Ardeshir Bastani, Gaurav Bajaj, Clifton Francis, Mahima Agumbe Suresh
San José State University

Abstract

Cloud computing offers an economical and elastic means to handle the storage and computation needs of the Internet of Things (IoT). However, storage and retrieval from the cloud could potentially violate policies, especially those pertaining to data privacy. Edge computing as a paradigm is a suitable way to overcome these issues. This poster presents an edge computing architecture that enables policy-aware normalization and filtration of the data that is sent to cloud services to preserve policies. We use a secure and encrypted channel to transmit the data generated by the IoT devices to the dedicated computing units at the edge of the network. Our architecture offers programmers the ability to configure the system easily and perform a predetermined set of computation tasks on the data, e.g., tasks to uphold privacy policies such as blurring faces, license plates, etc.

1 Introduction

The convergence of technologies such as wireless networks, commodity sensors, and embedded systems has led to the evolution of the Internet of Things (IoT). There are an increasing number of devices connecting to the Internet in applications including connected vehicles, home automation, wearable devices, connected healthcare, and appliances with remote monitoring capabilities. The exponential growth of IoT devices is sending overwhelming data volumes that need to be analyzed stored and protected [2].

Traditionally, data is transmitted to a cloud server for further processing. It is easier to maintain and secure the data stored in a centralized manner. However, there are significant challenges with this approach. Limited bandwidth and large volumes of data can significantly increase latency [3]. To overcome these challenges, cloud computing is experiencing a key move from the conventional style of acquiring centralized resources to distributed, decentralized architecture, called edge computing, where computing resources are placed closer to customers [4] [3]. Edge computing is an important paradigm to uphold policies concerning how the data is handled. E.g., A privacy policy may state that images containing faces and

licence plates will be blurred before making them available to third-party data processing tools.

This poster presents a proof-of-concept architecture that enables easy deployment of an edge cluster using Raspberry Pi devices running Kubernetes [1]. The architecture enables developers to set up the system using simple and easy configurations, which will need to specify what functions need to be executed and in what order. E.g., face detection followed by applying a blur function on a part of an image. Each data stream - processed and unprocessed will be available through Kafka queues. The developer will have to only configure access to the queue tokens. The architecture handles all of the other functionalities of the edge cluster including communication, on-boarding, logging, and resource management.

2 A Policy-Aware Architecture

In our architecture, IoT devices transmit their data streams, e.g., video data, using MQTT protocol to the collection of node servers acting as one unit of the device, called an *edge broker*. Edge broker is responsible for computing any data transmitted by the IoT devices. As soon as the computation of received data has been completed, it publishes results in a Kafka messaging queue. All the edge broker services will be running within a docker container which shall be orchestrated and managed by the Kubernetes cluster. Data is categorized into “public” and “private” types. By default, data is tagged as “private”. To impose policies on the data, a set of functions will be specified to be executed on the data. The resulting processed data may then be tagged as “public”.

References

- [1] Brendan Burns, Brian Grant, David Oppenheimer, Eric Brewer, and John Wilkes. Borg, omega, and kubernetes. *Queue*, 14(1):70–93, 2016.
- [2] L. Sanabria-Russo, D. Pubill, J. Serra, and C. Verikoukis. Iot data analytics as a network edge service. In *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pages 969–970, April 2019.
- [3] M. Satyanarayanan. The emergence of edge computing. *Computer*, 50(1):30–39, Jan 2017.
- [4] R. B. Wallace, H. Liu, R. Goubran, M. Bilodeau, and F. Knoefel. Streaming pressure sensor data for cloud processing. In *2019 IEEE International Symposium on Medical Measurements and Applications (MeMeA)*, pages 1–5, June 2019.