

IoT Data Analytics as a Network Edge Service

Luis Sanabria-Russo, David Pubill, Jordi Serra, Christos Verikoukis
Centre Tecnològic de Telecomunicacions de Catalunya (CTTC/CERCA)
{luis.sanabria, david.pubill, jordi.serra, cveri}@cttc.es

Abstract—Current IoT trends reveal an increase in computational requirements for data processing. Traditionally, data from sensors was uploaded to compute nodes at a backend cloud. Nevertheless, ever-growing amount of data generated by IoT devices have rendered this option too expensive in terms of network traffic, possibly leading to delays due to bottlenecks. Moreover, even if network connectivity were to be guaranteed, live processing of sensitive data (e.g.: biomedical) at a remote location may not comply with data protection policies. A popular approach tries to circumvent these issues by performing computational operations locally, that is, at the IoT Gateway level. This demo leverages open source lightweight virtualization tools and a container orchestration engine (i.e.: Docker and Kubernetes, respectively) in a cluster of IoT devices at the edge of the network, enabling the creation of a distributed pool of computing resources on top of which data analytics algorithms could be deployed, updated, or terminated. This approach guarantees that resource-hungry operations, such as live monitoring and real-time processing of sensitive data, are performed locally, reducing the overall delay and without risking data leaking to the outside world.

Index Terms—Virtualization, IoT, devOps, edge, container orchestration.

I. INTRODUCTION

Making sense of data generated by sensors requires both data analytics algorithms and a communication infrastructure that conveys the data from the sources to the computing nodes. Traditionally, sensor data was sent to a backend cloud, which depending on the volume may turn into a network bottleneck. The effects of this centralized approach adds congestion to the network in a south-north directions, provide a single point of failure, and further increases the monthly cost associated with the infrastructure (OPEX) if using an external provider, like Google Cloud, or AWS.

In order to circumvent the congestion issue, network architects follow a tiered IoT infrastructure topology which allow for some of the data processing to be performed at networking nodes, closer to where data are generated [1]. This intermediary intelligence enables pre-processing of the data, e.g.: caching, aggregation, reduction; which may significantly reduce the amount of traffic traversing the network towards the centralized storage. Figure 1 shows an example of a tiered architecture, where mid-level nodes or IoT gateways are equipped with processing and storage resources.

Despite being effective conceptually, manually maintaining such a distributed infrastructure is challenging. Redistribution of new code versions, and network re-configurations require access to each IoT gateway, hindering automation. Moreover, failure at a determined

IoT gateway may prevent it from capturing sensor data, disrupting the collection of possibly sensible data, as it is in the case with e-health applications.

To leverage the bottleneck, scalability, and reliability issues mentioned above, in this demo we implement a local cloud infrastructure based on open source software, as shown in Figure 1. It enables embedded intelligence at the level of IoT gateways, like in [2], but with the added reliability provided by a Kubernetes cluster for Docker containers [3], which is able to replicate software throughout the local cloud and ensure rolling updates that yield virtually zero downtime. Furthermore, to ensure security and reduction of south-north traffic, pre-processing and caching using persistent volumes are enabled at IoT gateways. Thereby only selected results from the pre-processing procedure are uploaded to centralized storage services, ensuring raw data are kept local.

An overview of the technological enablers for this demo is presented in Section II. Then, Section III provides a description of the proposed demo. Conclusions and future directions are contained in Section IV.

II. TECHNOLOGY ENABLERS

Collecting long time series of biomedical data is of special interest for e-health data analytics [4]. Specifically, taking care of elderly or ill people requires constant monitoring, not only of the patient’s activity, but also of its vital signs. Usually, long time series are fed to data models whose results are then used to estimate the well-being of a patient. In this light, the availability of communication, compute and storage infrastructure; as well as its ability to respond to events i.e.: alarm generations due to surges on determined metrics; are of utmost importance to active aging [5], and e-health applications.

In order to realize an architecture such as the one shown in Figure 1, a distributed set of IoT gateways must compose an edge cloud. This edge cloud should maintain persistent storage volumes for caching, as well as provide enough processing resources for algorithms performing analytics on the local data before their results could be uploaded to a backend cloud instance.

In this demo, all these requirements are fulfilled by building a Kubernetes (v1.10 or greater) cluster at IoT gateway level. Kubernetes is a container orchestrator engine, capable of scheduling the spawn and destruction of containerized applications (e.g.: Docker containers) throughout the cluster. Moreover, replicas and updates can be rolled-out without down-time, providing increased reliability.

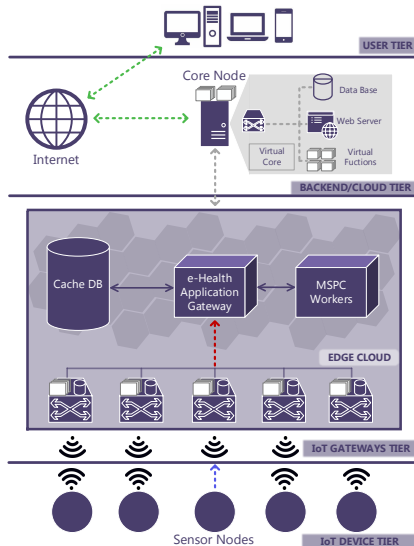


Fig. 1. An example end-to-end, tiered IoT platform architecture.

IoT gateways are central to the proposed architecture. Apart from composing the edge cloud, they could also work as receivers for sensor data by forming an ad-hoc WiFi network with sensor nodes. In this demo, Raspberry Pi 3 Model B devices are used as IoT gateways, while Raspberry Pi Zero W are used as sensor nodes. Sensors are attached to Raspberry Pi Zero W devices, which run a simple script that gathers metrics from sensors and transmits them to the closest IoT gateway using WiFi.

III. DEMO

This demo builds a cloud infrastructure which provides the required compute, storage and network resources for performing data analytics at the network edge. It instantiates a deployment emulating a health care facility, where IoT gateways are geographically distributed and connected among each other and the Internet using an L2/L3 wired network. On the other hand, patients are assumed to be equipped with wearable sensor node devices.

IoT gateways compose the edge cloud using Kubernetes over the Ethernet network, and employ WiFi interfaces to enable ad-hoc networking with sensor nodes. Patient-generated sensor metrics are emulated via software, and are transmitted via sensor nodes's WiFi interface towards the edge cloud.

Using Figure 1 as reference, when sensor data arrives at the edge cloud the e-Health Application Gateway classifies it according to the type of metric, patient ID, and time of day. Then, data is pushed to a persistent cache data base (Cache DB).

In order to detect deviations from each patient's healthy trend, a Multivariate Statistical Process Control (MSPC) per patient is periodically computed by a MSPC Worker using data from the Cache DB (via the e-Health Gateway). If the MSPC Worker finds a metric exceeding a predefined threshold, then the e-Health Gateway will be instructed to generate an alarm to

a health care professional in the form of a graphical plot. Additionally, selected MSPC Worker results (i.e. the graphical plot of the normalized variable values) are uploaded to a backend data base at the Core Node (see Figure 1) for storage and external visualization via Internet.

IV. CONCLUSION AND FUTURE WORK

The proposed IoT platform architecture allows for reliable reception, storage, and instantiation of one or several IoT data analytics algorithms at a local edge cloud. These functions are key for providing pervasive monitoring and alarm/event triggering for IoT applications, like elderly care. Data analysis tools such as Multivariate Statistical Process Control (MSPC), is used to detect deviations from trends in the incoming data.

This demo shows such architecture and application, generating a MSPC Worker per patient, and triggering an alarm via the generation of a plot if a determined metric exceeds a predefined threshold.

The proposed architecture can be further exploited in order to derive more accurate predictions. For instance, the MSPC workers can be replaced (or complemented) with more advanced distributed Machine Learning (ML) services, which interact among each other in order to handle bigger data sets and provide an additional level of data privacy [6]. Moreover, as IoT gateways are conceived using embedded devices such as the Raspberry Pi 3 Model B, support for other radio technologies (e.g.: Bluetooth, ZigBee, and others) is possible by means of USB adapters.

ACKNOWLEDGEMENTS

This work has been funded by the following research projects: SEMIOTICS (780315), SPOT5G (TEC2017-87456-P) and by the Generalitat de Catalunya under grant 2017 SGR 891.

REFERENCES

- [1] S. K. Sharma and X. Wang, "Live data analytics with collaborative edge and cloud processing in wireless iot networks," *IEEE Access*, vol. 5, pp. 4621–4635, April 2017.
- [2] K. Jaiswal, S. Sobhanayak, A. K. Turuk, S. L. Bibhudatta, B. K. Mohanta, and D. Jena, "An iot-cloud based smart healthcare monitoring system using container based virtual environment in edge device," in *2018 International Conference on Emerging Trends and Innovations In Engineering And Technological Research (ICETIETR)*, July 2018, pp. 1–7.
- [3] D. Pizzolli, G. Cossu, D. Santoro, L. Capra, C. Dupont, D. Charalampos, F. De Pellegrini, F. Antonelli, and S. Cretti, "Cloud4iot: A heterogeneous, distributed and autonomic cloud platform for the iot," in *2016 IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*, Dec 2016, pp. 476–479.
- [4] F. Firouzi, B. Farahani, M. Ibrahim, and K. Chakrabarty, "From EDA to IoT eHealth: Promise, Challenges, and Solutions," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, pp. 1–16, 2018.
- [5] I. Bisio, A. Delfino, F. Lavagetto, and A. Sciarrone, "Enabling IoT for In-Home Rehabilitation: Accelerometer Signals Classification Methods for Activity and Movement Recognition," *IEEE Internet of Things Journal*, vol. 4, no. 1, pp. 135–146, Feb 2017.
- [6] J. Konečný, H. B. McMahan, D. Ramage, and P. Richtárik, "Federated Optimization: Distributed Machine Learning for On-Device Intelligence," *CoRR*, vol. abs/1610.02527, 2016. [Online]. Available: <http://arxiv.org/abs/1610.02527>