Applying Security to a Big Stream Cloud Architecture for the Internet of Things

Laura Belli, University of Parma, Parma, Italy

Simone Cirani, University of Parma, Parma, Italy

Luca Davoli, University of Parma, Parma, Italy

Gianluigi Ferrari, University of Parma, Parma, Italy

Lorenzo Melegari, University of Parma, Parma, Italy

Marco Picone, University of Parma, Parma, Italy

ABSTRACT

The Internet of Things (IoT) is expected to interconnect billions (around 50 by 2020) of heterogeneous sensor/actuator-equipped devices denoted as "Smart Objects" (SOs), characterized by constrained resources in terms of memory, processing, and communication reliability. Several IoT applications have real-time and low-latency requirements and must rely on architectures specifically designed to manage gigantic streams of information (in terms of number of data sources and transmission data rate). We refer to "Big Stream" as the paradigm which best fits the selected IoT scenario, in contrast to the traditional "Big Data" concept, which does not consider real-time constraints. Moreover, there are many security concerns related to IoT devices and to the Cloud. In this paper, we analyze security aspects in a novel Cloud architecture for Big Stream applications, which efficiently handles Big Stream data through a Graph-based platform and delivers processed data to consumers, with low latency. The authors detail each module defined in the system architecture, describing all refinements required to make the platform able to secure large data streams. An experimentation is also conducted in order to evaluate the performance of the proposed architecture when integrating security mechanisms.

KEYWORDS

Big Stream, Cloud Computing, Internet of Things, OAuth, Open-Source Applications, Real-Time Applications, Security, Smart-X Applications

INTRODUCTION

In recent years, the forecast of a worldwide network of pervasively deployed heterogeneous networks is becoming a reality. The Internet of Things (IoT) involves billions of different devices, connected in an Internet-like structure, allowing new forms of interaction between things and people. The actors involved in IoT scenarios have extremely heterogeneous characteristics, in terms of processing and communication capabilities, energy supply and consumption, availability and mobility, spanning from Smart Objects (SOs) - i.e., constrained devices equipped with sensors or actuators, smartphones, wearables and other personal devices - to Internet hosts and the Cloud.

In order to allow heterogeneous nodes to efficiently communicate with each other and with existing Internet actors, shared and interoperable communication mechanisms and protocols are currently being defined and standardized. The most prominent driver for interoperability in the IoT is the adoption of the Internet Protocol (IP), namely IPv6. An IP-based IoT can extend and operate

DOI: 10.4018/IJDST.2016010103

with all existing Internet nodes, without additional efforts. Standardization institutions, such as the Internet Engineering Task Force (IETF), and several research projects are contributing to the definition of new mechanisms to bring IP to SOs (e.g., 6LoWPAN (Kim, Kaspar, & Vasseur, 2012)). This is motivated by the need to adapt higher-layer protocols (e.g., application-layer protocols) to constrained environments. As a result, IoT networks are expected to generate huge amounts of data, which can be subsequently processed and used to build several useful services for final users. The Cloud has become the natural collection environment for sensed data retrieved by IoT nodes, due to its scalability, robustness, and cost-effectiveness. Figure 1 shows the hierarchy of different levels involved in data collection, processing and distribution in a typical IoT scenario.

Sensed data are collected by SOs, deployed in IoT networks, and sent uplink to the Cloud which operates as collection entity and service provider. In some cases, intermediate Local Network Collectors (LNCs) can perform some preliminary tasks before sending data uplink, such as temporary data storage, data aggregation, and protocol translation. The described model is extremely general, and can be applied to different IoT scenarios. As an example, LNCs can be implemented by border routers or proxies.

Several relevant IoT application environments (e.g., industrial automation, transportation, and monitoring) require real-time performance guarantees or, at least, a predictable latency. Moreover, the performance requirements (e.g., in terms of data sources) may change even abruptly. On one hand, the large number of data sources represented by IoT nodes, generating a high rate of incoming data, and, on the other hand, the low-latency constraints call for innovative Cloud architectures able to handle efficiently such massive amounts of information.

A possible solution is given by Big Data approaches, developed in the last few years and become popular due to the evolution of online and social/crowd services, which are able to address the need to process extremely large amounts of heterogeneous data for various purposes. However, these techniques typically have an intrinsic inertia (as they are based on batch processing) and focus on the data itself, rather than providing real-time processing and dispatching (Zaslavsky, Perera, & Georgakopoulos, 2013; Leavitt, 2013). For this reason, Big Data approaches might not be the right solution to manage the dynamicity of IoT scenarios with real-time processing. In order to better fit

Figure 1. Different actors and layers involved in IoT scenarios: sensed data are sent from IoT networks to the Cloud, which provides services to consumers. At an intermediate level, preliminary local operations, such as data collection, processing, and distribution, may be carried out

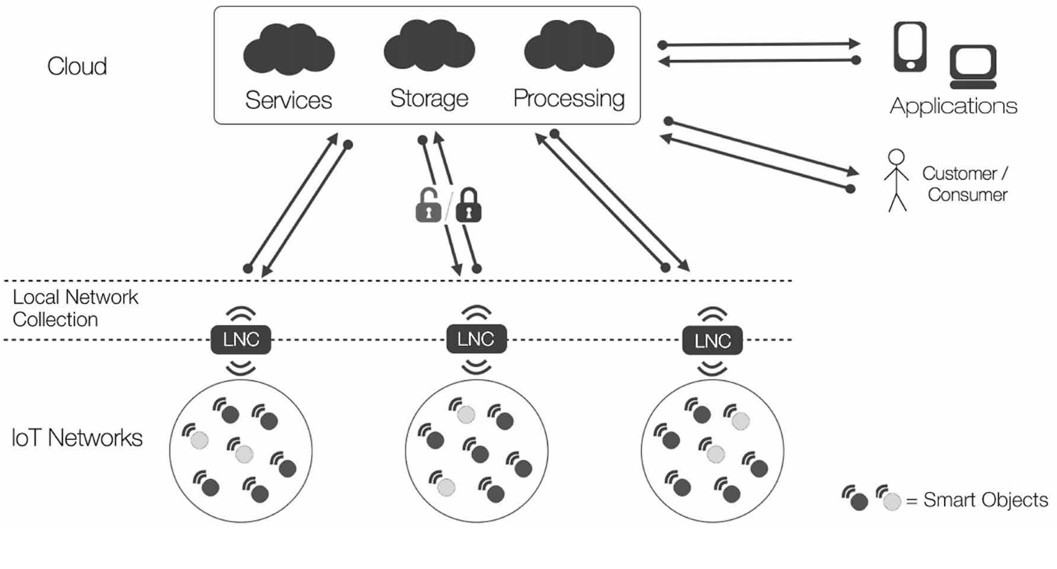
Cloud
Services
Storage
Processing
Applications
Customer / Consumer
Local Network Collection
LNC
LNC
LNC
IoT Networks
= Smart Objects