



## **The Art of Low-Cost IoT Solutions**

**13 June 2017**

**By Igor Ilunin, DataArt**

**[www.dataart.com](http://www.dataart.com)**

**© 2017 DataArt**

## Contents

Executive Summary.....	3
Introduction .....	3
The Experiment .....	3
The Setup .....	4
Analysis / Calculations .....	6
Pros, Cons and Considerations .....	7

## Executive Summary

- Global technology consultancy, DataArt set out to prove that it is possible to create a low-cost IoT solution.
- Currently delivering an IoT project has a high capital and operating expenditure.
- It is often a difficult choice for enterprises between the differencing tools and frameworks in the market – all of which have pros and cons.
- To prove that IoT does not have to be expensive, or difficult to establish DataArt selected flexible, reliable and cost-efficient platforms with minimal investment in infrastructure, software, knowledge and employees needed.
- The solution must also have to be proven to be scalable so that working firms could reasonably employ it.
- The below proof of concept could be used by start-ups to scale-ups that require their ideas to be validated quickly, at minimal cost.

## Introduction

We are living in the age of unprecedented industry disruption caused by a rapid speed of technological development, particularly in the area of the Internet of Things (IoT). IoT helps to transform industries into a data-driven paradigm, opening up vast opportunities. Fast IoT adoption is supported by companies which are transforming their business through technological innovation; manufacturers are providing low-cost and high-end devices and IoT platforms enabling device integration and management.

There are a lot of IoT tools and frameworks with different pros and cons present on the market, and it is a difficult choice to make. We believe that digital IoT adoption should turn towards flexible, reliable and cost-efficient platforms with minimal investment into infrastructure, software, knowledge and employees.

## The Experiment

To prove this concept, we decided to find out how to build an IoT solution from scratch with zero infrastructure and maintenance costs and little-to-no operating expense. This approach perfectly aligns with the cloud concepts such as pay for

what you use, pay as you go and Platform as a Service (PaaS), Software as a Service (SaaS), and serverless approaches.

## The Setup

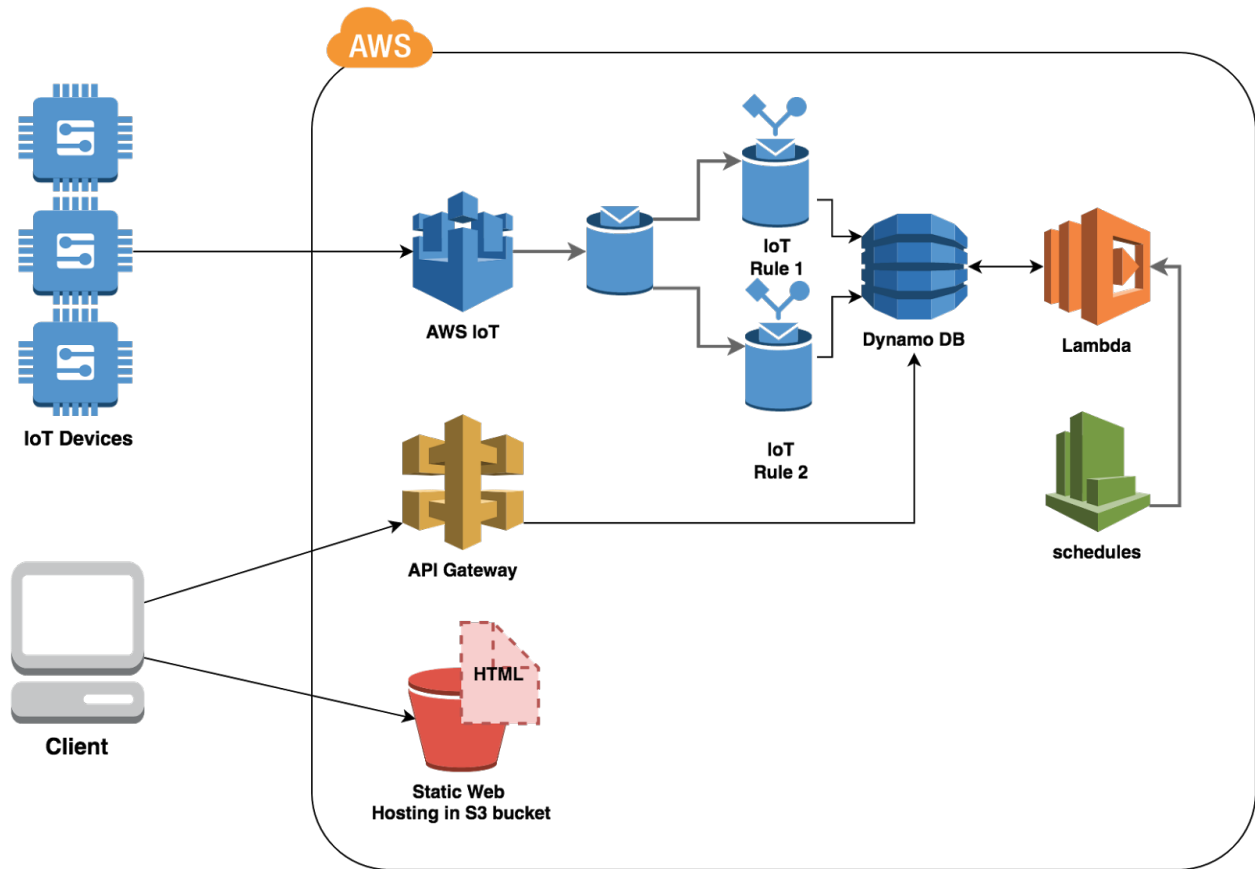
To implement the concept, we employed Amazon Web Services (AWS) cloud capabilities, as it is the most modern and innovative cloud on the market. Our project was to create device simulators reporting telemetry in real time and to have this information accessible in real time through the API.

We considered the AWS IoT platform as a powerful IoT framework. It supports MQTT - one of the most widely-used communication protocols. Our choice of supporting services to persist and process data was also based on its pricing and maintenance costs.

After researching a few different approaches to architecture and a variety of services, we finally came up with the following list of components:

- AWS IoT for data ingestions and device management,
- DynamoDB - document storage to persist data readings,
- AWS Lambda - serverless data processing,
- S3 - block storage used as static website hosting,
- Gateway API - REST access to the data.
- 

Fig i. Dataflow



The overall data flow for our concept project worked in the following way:

1. Devices are sending small volumes of data to AWS IoT (5 seconds each),
2. AWS IoT stores the data into DynamoDB table\*
3. Lambda functions are triggered each minute and each hour to do data analysis and store results back into DynamoDB,
4. API Gateway exposes data from DynamoDB to the world through the REST API,
5. Static HTML website is hosted on the S3, and is consuming REST API to display real-time data charts and analytics.

\*Point #2 may look a bit silly at first sight, as you could argue that DynamoDB is not the best choice to store raw time-series data. However, it does a great job for the demo purposes with a small number of devices.

We also considered using Firehose as a delivery stream from IoT to S3/Reshift and EMR cluster for data processing, but it would be an overkill for this micro-project.

We set the following key parameters for the architecture:

1. Free if you do not have devices reporting any data. Also with Amazon free tier, you can get it for free for the small number of devices.
2. Highly scalable and available out of the box due to the nature of each selected component.
3. Minimal knowledge required to start. You can just write the code, as you need only to define rules and write logic in one of the very popular languages: JavaScript, Python or Java.

## Analysis / Calculations

Let's say your backend operations are only required to process only several requests a minute – this means that most of the time your CPU is idle. We assume that you do not want to pay for the idle time. So here comes serverless architecture to ensure you do not.

Let's say you have 10,000 devices reporting small volumes of data every 15 minutes and this results in approximately 29.2 million requests monthly over 730 average hours a month.

AWS IoT costs \$5 per 1M requests, and DynamoDB costs \$0.0065 for every ten put requests per second and the same cost for every 50 get requests per second.

We get a monthly bill at around \$146 just to receive the data through AWS IoT, and \$14 to store it in DynamoDB with the smallest read capacity. We have \$160 in total which is equivalent to \$0.02 per device on a monthly basis or \$0.000005 per request. It does not take into account the usage of lambda, storage, API gateway which will actually be just a fraction of these numbers, so we can neglect it.

It is impressive to get an IoT solution with thousands of connected devices which will cost you less than two hundred dollars per month. However, let's imagine a business where devices are reporting critical data each second, and you have thousands of them. How much will you pay for serverless operations?

I will save you some time and just provide the results for 10k devices sending one message a second — the monthly payment will be more than \$136k and for 100k devices — \$1,36M accordingly and monthly cost per device increases to \$13.61. Would you still consider building your IoT solution serverless? The answer is - it depends!

## Pros, Cons and Considerations

All these figures mean that optimising your request rate will immediately and almost linearly lead to reducing your monthly bill.

This brings us to the second important conclusion we must consider — the total cost of ownership. There is a virtual threshold beyond which the serverless approach becomes very expensive and may not be effective.

For instance, the traditional architecture implementation cost may not depend so significantly on the number of devices or the number of requests per second, but the trade off is in additional operational expenses. The cost can also be reduced by using open source solutions.

Doubtlessly, the serverless architecture has many advantages:

- It transforms capital expenditure into operating expenditure, and generally reduces operational costs;
- You do not have to think about internal system administration processes;
- It reduces the development and deployment costs and timeframes (faster time to market);
- It is scalable and fault tolerant by design.

The first factor to consider is choosing the right approach for the requirements of your project. The serverless approach will be most appropriate if you are not concerned about the cloud lock-in and:

1. You are running a startup where you need validate your ideas fast, or you have a short time to market.
2. Your solution does not require frequent transmission of data from device to the cloud and therefore keeps the cost per device at a relatively low level.

On the other hand, you may consider custom or open source IoT solutions if you are building a cloud-agnostic and highly customisable solution and operating with real-time data.

For further information on DataArt's IoT development capability, please contact [igor.ilunin@dataart.com](mailto:igor.ilunin@dataart.com).

Note:

No part of this document may be reproduced or transmitted in any form without written permission from DataArt. The provision of the information contained herein is not intended to, and does not, grant any right or license under any DataArt patents or copyrights. DataArt shall have no responsibility to update this information.