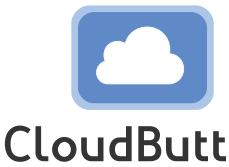

HORIZON 2020 FRAMEWORK PROGRAMME

CloudButton

(grant agreement No 825184)

Serverless Data Analytics Platform

D6.1 Communication Plan

Due date of deliverable: 30-06-2019
Actual submission date: 26-07-2019

Start date of project: 01-01-2019 Duration: 36 months

Summary of the document

Document Type	Report
Dissemination level	Public
State	v1.0
Number of pages	8
WP/Task related to this document	WP6 / T6.1, T6.2
WP/Task responsible	RHAT
Leader	Tristan Tarrant (RHAT)
Technical Manager	Marc Sánchez (URV)
Quality Manager	Gerard París (URV)
Author(s)	Tristan Tarrant (RHAT)
Partner(s) Contributing	RHAT
Document ID	CloudButton_D6.1_Public.pdf
Abstract	Overall communication plan and dissemination activities for the CloudButton project
Keywords	Communication, Dissemination, Conferences, Papers

History of changes

Version	Date	Author	Summary of changes
0.1	15-06-2019	Tristan Tarrant	First version
0.2	04-07-2019	Tristan Tarrant	Addressed reviewer's comments
1.0	26-07-2019	Tristan Tarrant	Final version

Table of Contents

List of Abbreviations and Acronyms

ACM	Association for Computing Machinery
BDV	Big Data Value
EMBL	European Molecular Biology Laboratory
FaaS	Function as a Service
LiDAR	Light detection and ranging
RHAT	Red Hat
TPI	The Pirbright Institute
URV	Universitat Rovira i Virgili
WP	Work Package

1 Executive summary

In order for the CloudButton project to successfully reach target groups, both in the short and in the long term, we need to deliver information in a clear, efficient and timely manner. The CloudButton Communication Plan introduces the strategy for dissemination, communication and exploitation, listing the guidelines that each partner, both industrial and research, should follow. The strategy will be tailored by each partner according to their specific involvement in the project and their respective target audience. The way that the project contents are communicated will naturally evolve over time: we foresee a gradual increase in the communication efforts as more results become available. Additionally, we need to take into account the fact that communication will continue after the formal end of the project. This document also includes the current external dissemination efforts made by the partners in the first part of the project.

2 Overall communication strategy

2.1 Key messages

The primary target in the communication strategy is to identify the key messages that the CloudButton project as a whole wants to deliver. These are:

- the definition and implementation of a highly scalable, high-performance serverless run-time architecture built around FaaS (Function as a Service) for Big Data analytics
- a mutable, in-memory, shared data layer delivering various data structures for serverless computing, which will offer persistence, coordination and concurrency control to the functions
- demonstration of the value and effectiveness of the above in the context of the analysis of large datasets in the field of bioinformatics (genomics and metabolomics) and geospatial data (LiDAR, satellite)

2.2 Target audiences

The CloudButton project aims to show how serverless computing, and in particular FaaS (Function as a service), is an extremely compelling paradigm for the deployment of heavy data processing applications. By leveraging the rising adoption of container- and microservice-based application architectures, developers have access to a strong foundation for the scalability requirements of these applications. The serverless platforms provide users with a simplified programming model, abstracting away most of the operational concerns while at the same time lowering the cost by optimizing resource usage to only those needed to run the chosen workload.

CloudButton will address the needs of developers by providing them with a zero-friction platform which will enable the quick development, deployment and management of both the computing and data layers. Built-in autoscaling capabilities will enable the automatic provisioning of additional capacity based on load and resource usage. These aspects will also benefit the researchers who will be able to focus solely on data processing and analysis without worrying about the underlying implementation details.

3 Dissemination and communication activities

We have identified the following communication channels as effective avenues for promoting the results of the CloudButton project:

- Academic and scientific conferences and publications
- Online communication
- Open-Source communities
- Customer engagements

3.1 Events

3.1.1 Conferences

Research results that have been achieved in the different work packages will be presented to the scientific community in the form of publications in journals, presentations and demonstrations in conferences, workshops and exhibitions. This is key to promote the ideas behind CloudButton, and to have the scientific backing for introduction in the relevant standards later on. The research quality of the academic partners and the good synergy with industrial ones should ensure high impact publications in top conferences and journals.

Some of the conferences where we are planning to attend are:

- Strata Data Conference (September 2019)
- KubeCon / CloudNativeCon Europe (May 2020)

3.1.2 Publications

We also plan to contribute to high impact academic events such as USENIX OSDI, ACM SOSP, USENIX NSDI, ACM EuroSys, IEEE ICDCS, ACM Middleware, USENIX ATC, ACM SuperComputing Conference, EuroPython, and PyCon in the second and third years of the project. We will also plan important demonstration activities targeted at the bioinformatics community and driven by EMBL and TPI in top events like ISCB, ECCB, NGS, and BioData World Congress among others.

These published results will follow a green open access strategy where articles and related data will be self-archived by the authors in an online repository (institutional repository, Zenodo, project's webpage) after the corresponding scientific publisher grants open access to authors. Most conferences and journals allow publishing the results after the peer-review period as long as the final copy-editing of the article is kept to subscribers or attendees of the journal and conferences respectively.

3.2 Online presence

The results of the CloudButton project are of a highly technical and specialized nature: therefore special care must be taken to present them to appropriate online communities (including social media, forums, video platforms) in a focused way in order to optimize the required effort.

3.2.1 Project website

The main channel of communication will be the project website (`http://cloudbutton.eu/`) which will offer both an overview of the project, the involved partners and the results that have been obtained. Additionally it will serve as an aggregation point for additional material (demos, quickstarts, videos) that will be delivered through external channels.

3.2.2 Social media

The primary social media channel is the official project Twitter account (@cloudbutton2020 `https://twitter.com/cloudbutton2020`) where announcements will be made about notable events, such as attendance at conferences, publishing of papers and availability of software.

Other social media channels will be used where appropriate. In particular we will use YouTube to deliver short simple, yet effective, demonstrations of the CloudButton design, architecture and implementation through the use of:

- "getting started", step-by-step examples
- self-contained quickstarts
- tutorials

3.2.3 Other online channels

Since most of the software deliverables of the project will be Open Source, it makes sense that they be made available through a collaborative software development platform.

The following lists the open source repositories where the project partners will be contributing:

- Crucial `https://github.com/danielBCN/crucial`
- Creson `https://github.com/otrack/creson/`
- Serverless Executor Service `https://github.com/otrack/serverless-executor-service/`
- Infinispan `https://github.com/infinispan/infinispan`
- Infinispan Kubernetes Operator `https://github.com/infinispan/infinispan-operator`

Additionally we intend to leverage online interactive learning platforms, like Katacoda (`https://www.katacoda.com/`) to offer users a chance to "try out" the CloudButton software without requiring them to install the entire software stack locally.

4 Community engagement

4.1 European and International initiatives

CloudButton is a member of the Big Data Value Public-Private Partnership (BDV). The BDV cooperation charter defines how members support its outreach and dissemination activities by using the association's branding and logo in the project's dissemination material, presenting results to the association itself as well as in connected EC events, such as participation at the European Big Data Value Forum and the Big Data Value PPP meet-up.

4.2 Open Source communities

Most of the software used and developed by the CloudButton project is open-source. It therefore makes sense to promote CloudButton to those communities which are closest in scope and requirements, including the Kubernetes and FaaS communities.

5 Customer engagements

The industrial partners in the project will need to define how they want to include the project results into their own marketing and sales pitches, tailoring the CloudButton message to their own requirements. Each partner will define how to use the results in the context of their existing portfolio. The following is a list of scenarios:

- as a self-standing product/service or as a new component of an existing product/service
- as a marketing tool to demonstrate the effectiveness of a product/service to interested parties

6 First-year activities

The following is a list of activities that have taken place since the beginning of the project:

- Online presence/Social Media
  - Post about the project in the Answare web page: `https://www.answare-tech.com/en/id-2/`
  - Dissemination of CloudButton and collaboration activities on Twitter (`https://twitter.com/metaspace2020`)
  - 06/02/2019 Project presentation post on The Matrix Foundation Website: `https://bit.ly/2XupqYt`
  - Added the CloudButton logo and link to the Infinispan homepage `https://infinispan.org`
- Conferences and Meetings
  - 27/02/2019 BDVA meeting in Brussels. Nuria de Lama, ARI representative at the BDVA Board, representing CloudButton and other Atos H2020 projects. Pedro García López presents CloudButton to BDVA community. Interactions with other H2020 projects: ELASTIC, SmartDataLake, and INFORE.
  - 01/03/2019 EMBL Presentation at the Highlight Symposium at European Molecular Imaging Meeting, Glasgow, UK.
  - 01/04/2019 Lighting talks at CS department, Technion. Dissemination about Big Data and CloudButton project
  - 21/05/2019 Stuttgart meetup "Serverless computing for data scientists and industry – your small step for a big difference" - Dr. Gil Vernik, IBM Haifa"
  - 13-15/06/2019 Pedro García (URV) attended JCSD 2019 (XXV Jornadas de Concurrencia y Sistemas Distribuidos) in Zaragoza, Spain, to disseminate preliminar results of the CloudButton project.
  - 26-28/06/2019 Red Hat and URV attendance at Big Data Value PPP Summit, Riga
- Papers
  - White-Box Atomic Multicast, A Gotsman, A. Lefort, G. Chockler, 49th IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'19)
  - State-Machine Replication for Planetary-Scale Systems, C. Baquero, V. Enes, T. França Rezende, A. Gotsman, M. Perrin, P. Sutra, submitted to SoCC'19.
  - On the FaaS Track: Building Stateful Distributed Applications with Serverless Architectures, D. Barcelona-Pons, M. Sánchez-Artigas, G. París, P. Sutra, P. Garcia Lopez, submitted to Middleware'19.
  - On the Correctness of Egalitarian Paxos, P. Sutra, submitted to Information Processing Letter.
  - ServerMix: tradeoffs and challenges of Serverless Computing, P. Garcia Lopez, M. Sanchez-Artigas, S. Shillaker, P. Pietzuch, D. Breitgand, G. Vernik, A. Juan, P. Sutra, work in progress.
- Other
  - 01/03/2019 ATOS Presentation of Cloudbutton project in ARI Newsletter
  - 06/03/2019 Interview at medium.com about Serverless Orchestration Services: `https://medium.com/@essenvanbas/f800564e8283`

- 13/03/2019 Radio Interview at Onda Cero to Pedro García López. Dissemination about Big Data and CloudButton project
- 26/03/2019 Radio Interview at Tarragona Ràdio to Pedro García López. Dissemination about Big Data and CloudButton project: `https://www.tarragonaradio.cat/contingut/urv_recerca_dimarts_26_de_marc_pedro_garcia_i_les_big_data/18039`
- 06/2019 Imperial remote presentation about parts of the CloudButton project to Intel Labs in Portland, USA.

7 Impact measurement

Measuring impact of dissemination activities is an important step in determining:

- whether the audience at the dissemination activity (conference, paper, meetup, social media, etc) was of the right kind for the project's subject matter
- whether the project's message was delivered correctly

During the duration of the project, the WP6 lead (Red Hat) will obtain quarterly information from the project partners about their dissemination activities, including both quantitative metrics as well as qualitative assessments. Such information will then be aggregated and redistributed to all the partners in a shared spreadsheet.

Examples of quantitative metrics are:

- attendance
- social media reach
- website analytics
- paper distribution
- newsletter subscribers

Qualitative assessments will vary depending on the channel, and will be used to determine whether any corrective measures should be taken in shaping the strategy of the communication strategy, e.g. whether to focus more on certain areas which have sparked interest.

8 Conclusions and next steps

To capture the attention of potential users of the CloudButton architecture, it is necessary to have a "demo-worthy" base that can be used to build basic examples and tutorials. We strongly believe in a hands-on approach which users can reproduce locally with ease. Past experiences have shown that users will accept some shortcomings and incompleteness in favour of early proof of the soundness of the approach. We shall therefore need to ensure that the results achieved during the course of the project are usable as early as possible and that incremental progress is highlighted regularly.