Project Title	AN OPEN, TRUSTED FOG COMPUTING PLATFORM FACILITATING THE DEPLOYMENT, ORCHESTRATION AND MANAGEMENT OF SCALABLE, HETEROGENEOUS AND SECURE IOT SERVICES AND CROSS-CLOUD APPS
Project Acronym	RAINBOW
Grant Agreement No	871403
Instrument	Research and Innovation action
Call / Topic	H2020-ICT-2019-2020 / Cloud Computing
Start Date of Project	01/01/2020
Duration of Project	36 months

D7.10 – Project Impact Assessment

Work Package	WP7 – Dissemination, Exploitation and Communication
Lead Author (Org)	Stefanos Venios (Suite5)
Contributing Author(s) (Org)	Chryssovalanto Kousseti (Suite5); Thomas Pusztai (TUW); George Pallis, Marios Dikaiakos, Demetris Trihinas, Moysis Symeonidis, Demetra Katziani (UCY); Heini Bergson Debes (DTU); Christina Stratigaki, Alexandros Vasileiou, Konstantinos Theodosiou (UBITECH); Ronald Toegl (INFINEON);Vasileios Psomiadis (AUTH); George Amponis (K3Y)
Due Date	31.12.2022
Actual Date of Submission	30.12.2022
Version	1.0

Dissemination Level

X	PU: Public (*on-line platform)
	PP: Restricted to other programme participants (including the Commission)
	RE: Restricted to a group specified by the consortium (including the Commission)
	CO: Confidential, only for members of the consortium (including the Commission)

The work described in this document has been conducted within the project RAINBOW. This project has received funding from the European Union's Horizon 2020 (H2020) research and innovation programme under the Grant Agreement no 871403. This document does not represent the opinion of the European Union, and the European Union is not responsible for any use that might be made of such content.

Versioning and contribution history

Version	Date	Author	Notes
0.10	31.03.2021	Stefanos Venios (Suite5)	Defined initial ToC
0.20	15.11.2021	Stefanos Venios (Suite5)	Updated ToC
0.30	09.11.2022	Chryssovalanto Kousseti (Suite5) Thomas Pusztai (TUW) George Pallis, Marios Dikaiakos, Moysis Symeonidis (UCY) Heini Bergson Debes (DTU) Christina Stratigaki, Alexandros Vasileiou, Konstantinos Theodosiou (UBITECH)	Input in sections 4, 5.1, 5.4 and 6.2
0.40	12.12.2022	Chryssovalanto Kousseti (Suite5) Demetris Trihinas (UCY) Christina Stratigaki, Alexandros Vasileiou, Konstantinos Theodosiou (UBITECH)	Input in sections 1, 2, 5.2 and 6.4 Updates in section 4
0.50	14.12.2022	Stefanos Venios (Suite5) Ronald Toegl (INFINEON) Christina Stratigaki, Alexandros Vasileiou, Konstantinos Theodosiou (UBITECH) Thomas Pusztai (TUW)	Input in sections 3, 5.1, 5.3, 5.4 and 6.1 Updates in sections 2.2 and 5.1
0.60	16.12.2022	Stefanos Venios (Suite5) Vasileios Psomiadis (AUTH) Demetra Katziani (UCY)	Input in Executive Summary and sections 6.2, 6.3 and 7
0.70	19.12.2022	Stefanos Venios (Suite5) George Amponis (K3Y)	Input in section 4.1 Submitted for internal peer review
0.80	22.12.2022	Stefanos Venios (Suite5)	Implementing internal peer review comments Updates in section 7
0.90	27.12.2022	Stefanos Venios (Suite5) George Amponis (K3Y) Thomas Pusztai (TUW) Christina Stratigaki (UBI)	Implementing internal peer review comments Updates in sections 4.1, 4.2, 5.3, 5.4 and 7
1.0	30.12.2022	Christina Stratigaki (UBI)	QA review and Submission

Disclaimer

This document contains material and information that is proprietary and confidential to the RAINBOW Consortium and may not be copied, reproduced or modified in whole or in part for any purpose without the prior written consent of the RAINBOW Consortium

Copyright © Rainbow Consortium Partners 2020 -2022

Despite the material and information contained in this document is considered to be precise and accurate, neither the Project Coordinator, nor any partner of the RAINBOW Consortium nor any individual acting on behalf of any of the partners of the RAINBOW Consortium make any warranty or representation whatsoever, express or implied, with respect to the use of the material, information, method or process disclosed in this document, including merchantability and fitness for a particular purpose or that such use does not infringe or interfere with privately owned rights.

In addition, neither the Project Coordinator, nor any partner of the RAINBOW Consortium nor any individual acting on behalf of any of the partners of the RAINBOW Consortium shall be liable for any direct, indirect or consequential loss, damage, claim or expense arising out of or in connection with any information, material, advice, inaccuracy or omission contained in this document.

Copyright © Rainbow Consortium Partners 2020 -2022

Table of Contents

Executive Summary *6*

1 Introduction *8*

1.1 Scope and Objectives of the Deliverable **8**

1.2 Structure of the Deliverable **8**

2 RAINBOW Research Impact Expectations *9*

2.1 Expected Impact **9**

2.2 Impact Assessment Approach **9**

3 Contribution to Strategic Impacts of Objective ICT-15-2019 ..... *12*

4 RAINBOW's Research Outputs *16*

4.1 Impact on Demonstrators **19**

4.1.1 Human-Robot Collaboration in Industrial Ecosystems 19
4.1.2 Digital Transformation of Urban Mobility 20
4.1.3 Power Line Surveillance via Swarm of Drones 22

4.2 Lessons Learnt **24**

5 Research, Policy, and Industry Community Benefits *27*

5.1 Contribution to Standardisation **27**

5.1.1 Contribution to TCG 27
5.1.2 Contribution to Fog Adoption IFG 27
5.1.3 Contribution to the Open Horizon initiative 28
5.1.4 Contribution to the Linux Foundation's Centaurus project 28

5.2 Alignment with OpenFog architecture **28**

5.3 Improving Innovation Capacity **31**

5.4 Contribution to European Policies and Other Strategic Initiatives **31**

5.4.1 RAINBOW Fog Computing Platform - HRP Entry 32
5.4.2 Fogify: A Fog Computing Emulator Framework – HRP Entry 34

6 Knowledge Transfer *37*

6.1 Contribution to TPM **37**

6.2 Dissemination, Communication, Exploitation **37**

6.2.1 Scientific Publications 41
6.2.2 Fogify 42

6.3 Hackathons **44**

6.4 RAINBOW Adoption Guidelines **45**

6.4.1 RAINBOW Installation 45
6.4.2 RAINBOW Usage Guide 49

7 Conclusions *57*

Copyright © Rainbow Consortium Partners 2020 -2022

List of tables

Table 4-1 Impact KPIs linked to the project's research output 16
Table 6-1 RAINBOW Scientific Publications 41
Table 6-2 .env variables of Analytics Stack - master node 47
Table 6-3 .env variables of Analytics Stack – worker nodes 48

List of figures

Figure 2-1 The Three Horizons Framework 10
Figure 2-2 RAINBOW impact assessment approach 11
Figure 5-1 OpenFog Reference Architecture and Perspectives 29
Figure 5-2 RAINBOW and OpenFog superposition 30
Figure 5-3 Horizon Results Platform – RAINBOW results 31
Figure 5-4 HRP - RAINBOW Fog Computing Platform-Result Title, Target Audiences and Needs 32
Figure 5-5 HRP - RAINBOW Fog Computing Platform-About us 32
Figure 5-6 HRP - RAINBOW Fog Computing Platform-Result Description and Influence 33
Figure 5-7 HRP - RAINBOW Fog Computing Platform-Result and Business Maturity and Exploitation outlook 33
Figure 5-8 HRP - RAINBOW Fog Computing Platform-Investors' Corner 33
Figure 5-9 HRP - Fogify: A Fog Computing Emulator Framework - Result Title, Target Audiences and Needs 34
Figure 5-10 HRP - Fogify: A Fog Computing Emulator Framework – About Us 34
Figure 5-11 HRP - Fogify: A Fog Computing Emulator Framework – Result Description and Influence 35
Figure 5-12 HRP - Fogify: A Fog Computing Emulator Framework - Result and Business Maturity and Exploitation outlook 35
Figure 5-13 HRP - Fogify: A Fog Computing Emulator Framework – Investors Corner . 36
Figure 6-1 Fogify architecture 43
Figure 6-2 Core platform installation scripts 46
Figure 6-3 Analytics Stack setup – master node 47
Figure 6-4 Analytics Stack setup – worker node 48
Figure 6-5 Dashboard component scripts 49
Figure 6-6 Main page of RAINBOW Dashboard 50
Figure 6-7 Components List in the RAINBOW Dashboard 50
Figure 6-8 Component Configuration in the RAINBOW Dashboard 51
Figure 6-9 The Service Graph Editor of RAINBOW Dashboard 51
Figure 6-10 Resources management in the RAINBOW Dashboard 52
Figure 6-11 Application deployment procedure in the RAINBOW Dashboard 52
Figure 6-12 Application Instances list in the RAINBOW Dashboard 53
Figure 6-13 Analytics and SLO Editor in the RAINBOW Dashboard 53

Copyright © Rainbow Consortium Partners 2020 -2022

Figure 6-14 Creating analytics metric in the RAINBOW Dashboard 54
Figure 6-15 Adding expressions to the analytics metric in the RAINBOW Dashboard.... 54
Figure 6-16 SLO creation in the RAINBOW Dashboard 55
Figure 6-17 Adding metrics to the SLO in the RAINBOW Dashboard 55
Figure 6-18 Adding computations to the SLO in the RAINBOW Dashboard 56
Figure 6-19 Adding Expressions to the SLO in the RAINBOW Dashboard..... 56

Copyright © Rainbow Consortium Partners 2020 -2022

Executive Summary

RAINBOW's impact assessment is our effort to prove the applicability, usability, effectiveness and value of the RAINBOW integrated framework in industrial, real-life trustworthy services, applications and scenarios, demonstrating the RAINBOW artefacts, methodologies and services under pragmatic conditions against the project's use cases.
The present deliverable concludes the work in Work Package 7 "Dissemination, Exploitation and Communication" and by extension, the whole RAINBOW project. Collecting input from almost all Work Packages of the project and from various partners, it highlights the most prominent achievements of RAINBOW project during the 36 months between 2020 and 2022.

Initially, it lays down the impact assessment approach, so as to define and employ a pertinent technique to comprehend RAINBOW's realized impact. Inspired by an extremely versatile strategic framework, called "The Three Horizons Methodology", we look for specific impact along three process stages: "research output", "knowledge transfer" and "community benefit".

Furthermore, starting from the envisioned impact the project would have towards the strategic goals of the Objective ICT-15-2019, as described in the DoA [1], we discuss the means through which this impact was achieved, and to which extend. Objective ICT-15-2019[1] called for "*competitive cloud solutions based on advanced cloud platforms, services, cloud-based software and data applications. Furthermore, such solutions should also address advanced security, data protection, performance, resilience and energy-efficiency requirements to respond to the future digitisation needs of industry and the public sector*".

Getting started with the impact assessment stages, we turn our attention towards the results of the project's technical work, in a way that sheds light on the impact this work has on the strategic objectives of the call. In parallel to that, we sketch the added value and present the advantages and disadvantages the developed solution has on three advanced fog use cases, i.e., the three demonstrators described in the project.

The focus is then turned towards certain community benefits, from contributions to standardisation initiatives through the provision of open source code (to the Eclipse IoT community and to Linux Foundation's Centaurus project), to the open sourcing of RAINBOW services (the RAINBOW orchestrator, the monitoring agent, the analytic enabler, the storage agent, the overlay mesh networking stack, etc.), to the alignment with the architectural principles of OpenFog Consortium (rebranded as the Industry IoT Consortium in August 2021).

As a last means to measure impact, we document the efforts to transfer the produced know-how to interested stakeholders. Such efforts include Webinars, Scientific Publications, Conference Papers, etc. in the context of dissemination efforts; website,

[1] https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/ict-15-2019-2020

Copyright © Rainbow Consortium Partners 2020 -2022

social media accounts, flyers, etc. as communication material; organisation of Hackathon events in partner universities, creation of user manual and adoption guidelines.

Copyright © Rainbow Consortium Partners 2020 -2022

1 Introduction

1.1 Scope and Objectives of the Deliverable

This deliverable's objective is to undertake a preliminary analysis of the impact generated throughout the project, based on a framework that categorizes the various outcomes because of the various research process stages. Research output, knowledge transfer, and community benefits are the stages of the research process, and they are represented by Horizons 1-3 in the Three Horizons framework. The best way to comprehend RAINBOW's realized impact is to look for specific impact along those three stages.

To achieve this, the deliverable at hand reports on the qualitative and quantitative assessment of project activities. First, it highlights the achievements in the context of produced research output. Also, it demonstrates the applicability, usability, effectiveness, and value of the RAINBOW integrated framework in real-world applications. Then, it collects the standardisation initiatives and alignment with existing frameworks. Finally, it reports on knowledge transfer efforts, in the dissemination and communication front, but also at the level of expanding into the student community and producing solid project adoption guidelines for future uptake of the achieved results.

1.2 Structure of the Deliverable

The deliverable is structured as follows:
Chapter 2 lays down the project's expected impact and draws the impact assessment approach.
Chapter 3 presents the contribution towards the strategic impacts identified in the related ICT-15-2019 call.
Chapter 4 highlights the project's technical work in a way that sheds light on the impact this work has on the project's stakeholders and demonstrators.
Chapter 5 presents the benefits of the project towards Research, Policy, and Industry Community.
Chapter 0 documents at high-level the project's Knowledge Transfer efforts.
Finally, **Chapter 7** derives the conclusions of the deliverable.

Copyright © Rainbow Consortium Partners 2020 -2022

2 RAINBOW Research Impact Expectations

2.1 Expected Impact

The ubiquity of smart devices and rapid development of modern virtualization and cloud technologies have brought fog/edge computing on the foreground, defining a **New Era in Cloud Computing**. In the same way that in the last decade the whole industry has moved from the traditional centralised data processing to the cloud, it is widely expected that **fog/edge computing is the future**, initially at corporate level and probably very soon at home/personal level as well.

RAINBOW "***supports the development of complete solutions which, enabled by cloud technologies, are able to take advantage of distributed infrastructure including local scaled-down network nodes, micro datacenters and connected devices***". The high-level challenge that the RAINBOW tackles and the expected impact of the project from this scope is clear: improving application performance and security without incurring the cost and inflexibility of local processing.

RAINBOW seeks to propose a new development paradigm capable of spawning an ecosystem that supports the development of future IoT services and cross-cloud applications. Thus, the basic concept is to allow developers, service providers, and infrastructure to focus on what they do best rather than developing complex services for smart resource allocation, monitoring, and service orchestration. Simultaneously, developers and IT solution providers (particularly SMEs that cannot afford to invest in expensive proprietary modules) can obtain a complete solution that ensures device and mesh network security, data protection, identity management, anonymity, and resource integrity at all networking layers. In terms of data management, RAINBOW addresses the complex issue of safely storing, processing, and sharing data at the fog and edge layer; not only will it provide real-time, interoperable, and secure data exchange resembling a "Decentralized Edge API" through which any IoT service, edge device, and cooperative fog node can share, access, and write data, but it will also integrate data-analytics and processing solutions via extensions scheduling interfaces of popular open-source frameworks.

RAINBOW has identified several groups towards which its final solution is targeted, which include mainly ICT SMEs, individual developers and researchers, early adopters, mainstream users, as well as the industrial and the governmental sectors.

2.2 Impact Assessment Approach

According to the RAINBOW DoA [1], the proposed way to comprehend RAINBOW's realized impact is to look for specific impact along three process stages "research output", "knowledge transfer" and "community benefit". Each of these research process stages includes specific outputs and outcomes that have been considered within a preliminary framework to properly assess the impact achieved within the project.

Copyright © Rainbow Consortium Partners 2020 -2022

To achieve this, the Three Horizons of Growth framework can be considered, an extremely versatile strategic framework, applicable to most organizations. The Three Horizons Methodology links the present with the desired future and identifies the "seen" disruptions that can occur on the way to a vision. This strategic framework requires that the objectives be divided into 3 different "horizons" (see Figure 2-1).

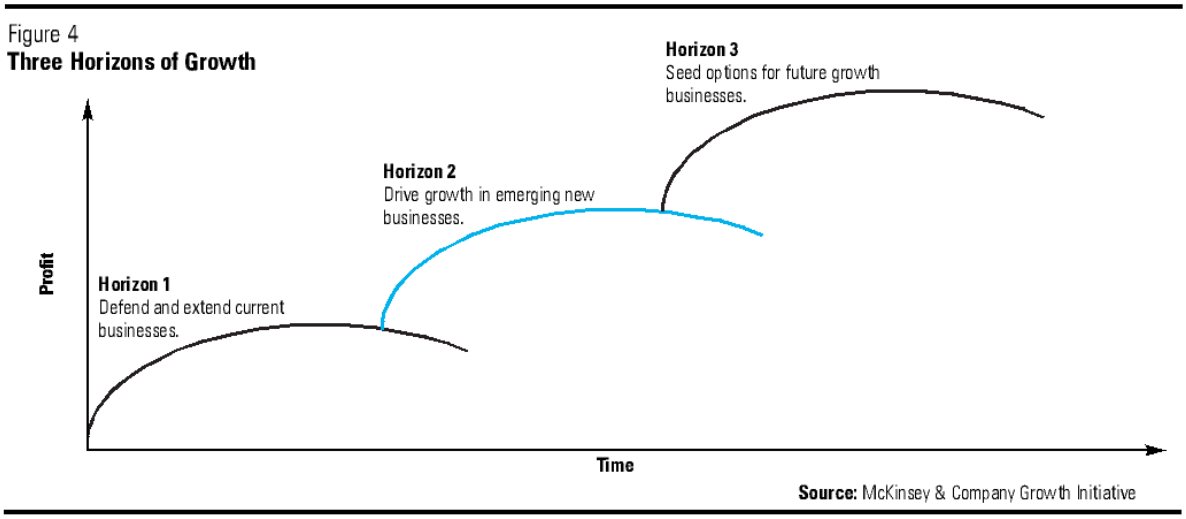

Figure 2-1 The Three Horizons Framework

Horizon 1 represents those core research domains most readily identified within the research project. Here the focus is on generating results to maximize the value. Horizon 2 encompasses emerging opportunities, including rising entrepreneurial activities likely to generate substantial profits in the future but that could require considerable investment. Horizon 3 contains ideas for profitable growth down the road—for instance, exploitation approaches, pilot programs, etc. This reflects a holistic approach towards impact creation. By mapping the logic model to impact creation processes, it is possible to draw a sequence of stages that lead from research to impacts.

The first stage **"research output"** is devoted to primary research outputs that emerge in the short or medium term as a direct outcome of the research endeavor. As a result, this research output can be assigned to concrete assets generated because of RAINBOW. The applicability, usability, effectiveness, and value of those RAINBOW results in real-world applications, i.e., the three demonstrators of the project, are also included in this stage. Building on the research output, a respectable number of scientific papers have been published and knowledge transfer events have been held (e.g. Hackathon), allowing the knowledge and conclusions gained in the first stage to be disseminated further in the second stage **"knowledge transfer"**. At this stage, in particular dissemination and communication must be considered together with the organised Hackathons (see section 6.3) and the exceptional research outputs achieved (see section 6.2.2), as they allowed to (a) validate the RAINBOW assets in the frame of use cases; (b) transfer the knowledge and the results to various user groups receiving valuable feedback. Finally, the third stage

Copyright © Rainbow Consortium Partners 2020 -2022

"community benefit" captures social, economic, or environmental impact that will occur over a longer period on the base of the previous RAINBOW results. This stage includes activities which provide policy related impact, contribution to standardisation groups, policy for the innovation capacity and project results' sustainability, as well as the RAINBOW adoption guidelines (see section 6.4), an enabler of future services based on RAINBOW.

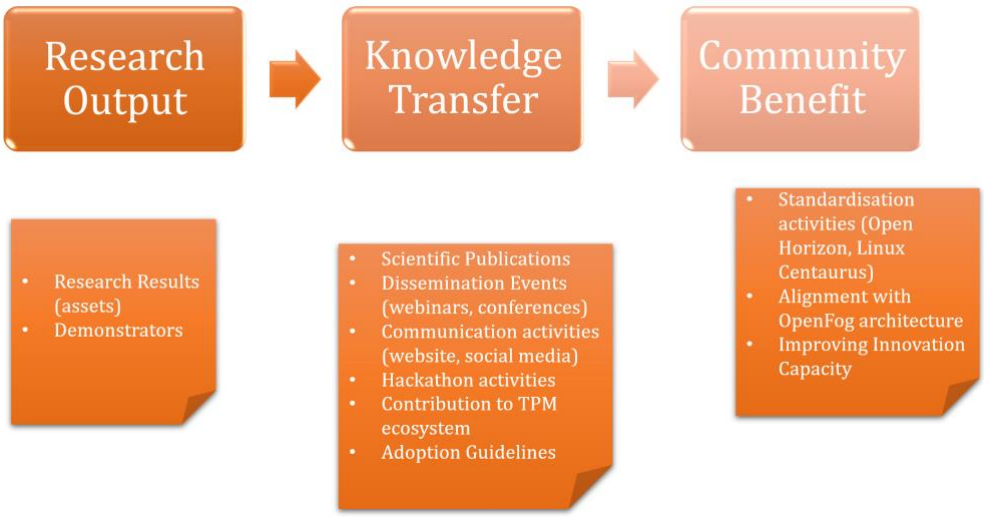

Figure 2-2 RAINBOW impact assessment approach

Copyright © Rainbow Consortium Partners 2020 -2022

3 Contribution to Strategic Impacts of Objective ICT-15-2019

With regards to the Horizon 2020 Work Programme 2018-2020, the pillar "Leadership in Enabling and Industrial Technologies" and in particular the Work Programme part "Information and Communication Technologies", RAINBOW contributes to the underpinning purpose of "maintaining and developing the technology leading edge in key areas" such as the one of cloud computing and for "digitizing and transforming European industry and services".

ICT-15 Strategic Impact #1 - *Contribute to the development of an ecosystem that will respond to the future digitisation needs of industry and the public sector*

RAINBOW aims at extending the Cloud towards the IoT to achieve improved QoS and empower latency-sensitive and bandwidth-hungry applications. RAINBOW responds to the actual needs of both the majority of the industrial domains and the public sector services as far as it concerns their digital transformation (enabling complement to the IoT & Edge and to the IoT & Cloud scenarios). The RAINBOW integrated platform (as described in deliverables D5.2 [2], D5.3 [3] and D5.4 [4]) has been designed and developed to achieve the following goals:

- Decrease in effort and investments for developing and managing the lifecycle of fog computing services and increase in software delivery cycles speed (compared to other state-of-the-art approaches);
- Increase in productivity of business applications, which rely on, or can be developed based on fog computing services (compared to the current status);
- Lowering access barriers for SMEs for the usage of advanced technologies for the development, management, security, and orchestration of services, made available to SMEs.

The integrated RAINBOW platform tackles all these issues by combining its state-of-the-art sub-components. Specifically:

- Starting from the automated on-boarding of the edge and cloud devices the RAINBOW platform alleviates the time-consuming creation of edge-to-cloud network fabric and guarantees secured communication via its Secure Overlay Mesh Networking as described in deliverable D2.6 [5]. The novel Mesh Routing protocol stack is realized via a library, which is responsible for the secure onboarding and operation of a consistent network overlay among the fog nodes, and the selection of a cluster-head.
- With the cluster of devices ready, RAINBOW users can utilize RAINBOW's intuitive dashboard in order to configure, deploy, and manage their applications. The RAINBOW dashboard and its usability are highlighted in deliverable D5.3 [3] and, specifically, in Section 5 which provides a usage guide of the platform. To this end, the RAINBOW user's interaction with the underlying infrastructure goes through visualizations and the user interface (UI) of the RAINBOW dashboard without

Copyright © Rainbow Consortium Partners 2020 -2022

users dealing with any time-consuming and manual configuration to the underlying cluster.

- Furthermore, users describe services and their settings via a Service Graph, and the underlying orchestration strategies handle cluster federation, fully SLO-aware scheduling, and additional user-defined elasticity strategies, as we described in deliverable D3.2 [6]. The latter functionalities improve productivity, minimize latency, and increase the performance of the deployed application.
- In addition, RAINBOW does not only provide the required infrastructure and application-level monitoring metrics [6], but also handles monitoring data management by providing novel distributed data storage and sharing services (Section 3 of [7]). Furthermore, RAINBOW processing and analytic services extract and process efficiently (based on trade-offs among performance, energy, and data quality) the monitoring metrics (Section 4 of [7]) sending them to the RAINBOW's orchestrator to perform actions (e.g., scaling) and visually depicting them on the RAINBOW dashboard derived from high-level analytic queries (Section 5 of [7]).
- Finally, we have to note that open-source services (a complete list of the open-source components found in [4]) that are created during the platform implementation, such as Fogify emulation framework (Section 4.2 of [7]) or RAINBOW's monitoring library [6], fulfil various needs (e.g., the need for effortless performance evaluation) of the Fog-Cloud industries and institutions, while standardization efforts, as they are described in [8], contribute in guidance of future digitization needs of industry and the public sector.

ICT-15 Strategic Impact #2 - *Assist the development of new cloud-based services and infrastructures in Europe and foster an industrial capability in the cloud computing sector*

The trusted fog computing platform which is the main outcome of RAINBOW facilitates the development of heterogeneous, secure, and privacy-preserving services, incorporating technological and business requirements directly coming from the industry [9], [10], [11] and our early implementation of our demonstrator applications (Human-Robot Collaboration, Digital Transformation of Urban Mobility, Power Line Surveillance early releases). Through the evaluation of RAINBOW use cases, we highlight the benefits of RAINBOW, and showcase the deployment procedure of existing applications on it. Specifically, the Human-Robot Collaboration Demonstrator in [12], Digital Transformation of Urban Mobility demonstrator in [13], and Power Line Surveillance Demonstrator in [14] showcased the deployment pipeline and proved the positive KPIs impact that the platform provided to them in [15] [16]. RAINBOW platform achieved to simplify significant crucial and complex industrial processes, e.g., device management, deployment via RAINBOW dashboard [4], etc., while enhancing applications' features with increased security via security enablers and mesh routing protocol stack [5], improved performance via orchestration and auto-scaling management [6], decreasing energy consumption (via adaptive monitoring and analytic processing scheduling that are highlighted in [6] [7]), and so on. Furthermore, our demonstrators validated the platform's features in the [15] [16] in the before-mentioned

Copyright © Rainbow Consortium Partners 2020 -2022

axis. The final version of RAINBOW, as described in [4], gives a new prospect for the development of innovative Fog and cross-cloud services directly exploitable at the industrial level by facilitating the deployment and management of scalable, heterogeneous, and secure IoT services and cross-cloud applications (i.e., microservices).

ICT-15 Strategic Impact #3 - *Create new opportunities to encourage European-based providers, in particular SMEs, to develop and offer cloud-based services based on the most advanced technologies*

IT providers are able to create their own services on top of the RAINBOW via:

- Communicating and interacting through a wide range of well-defined and open APIs like they are described in deliverables D5.2 [2], D5.3 [3] and D5.4 [4] and in separate components' deliverables [6] [7] [5] along with components configuration properties. We should note here, there is also a more technical description of components interaction in the documentation site (https://rainbow-h2020.readthedocs.io) and in its component repository (more details about repositories can be found in each component's respective deliverable).
- Furthermore, many RAINBOW services are open source, e.g., the RAINBOW orchestrator, the monitoring agent, the analytic enabler, the storage agent, the overlay mesh networking stack, etc. as they described also in Deliverables [17] [11]. SMEs and technological industries can encapsulate these components in their software stack by integrating them via their API interfaces.
- Many RAINBOW components provide out-of-the-box integration with well-known cloud software tools. For instance, the orchestration algorithms and the secured RAINBOW overlay mesh networking are already integrated with the most well-known cloud-based orchestrator, namely Kubernetes [6] [5], and the RAINBOW Fog Analytics Service can automatically translate the users' high-level queries to various big data streaming processing engines like Apache Spark and Storm described in [7].

To this end, RAINBOW components expose open and easy-to-use APIs while also trying to enhance popular and well-known software stacks via APIs, direct integration, or code extension to provide more advanced technologies and be acceptable to the industry. Last but not least, the standardization efforts in [8] would influence Fog computing architectures of future applications and allow them to benefit from the outcomes of the RAINBOW project.

ICT-15 Strategic Impact #4 - *Leverage research and innovation projects to support the development and deployment of innovative cloud-based services and next generation applications, for the public and private sectors (including standardisation and applications for Big-Data and other sector specific applications)*

The RAINBOW integrated platform is built on top of innovative software solutions and middleware tools guaranteeing a secure and near-to-optimal deployment. Specifically:

Copyright © Rainbow Consortium Partners 2020 -2022

- a novel data collection process with adaptive monitoring and adaptive dissemination of captured metrics [6],
- innovative Fog orchestration mechanisms with also real-time adjustments [6],
- enriched security with attestation security enablers and secured data exchange via mesh networking [18] [5],
- a highly available data storage with smart replication mechanisms and fast monitoring data retrieval APIs [7],
- adaptable monitoring and dissemination algorithms for the minimization of the monitoring energy consumption and data transfer [6],
- declarative analytics and efficient real-time execution of them with configurable placement algorithms [7].

We validated the latter contributions and advantages of the RAINBOW platform through:

- The implementation and examination of RAINBOW's use cases as described in [12] [13] [14] and evaluated in [15] [16], which are characteristic examples of next-generation internet (NGI) and IoT applications.
- The scientific outcomes of the project, with more than 20 publications in top tier conferences and journals and funded 6 PhD students' theses contribute to further research in the tackled domains and for new research and innovation projects on the development of real-life fog-edge-cloud-based services and applications. A more detailed overview of the RAINBOW's contributions and exploitation can be found in [17] [19] [11].
- Moreover, sub-components of the project are already open-sourced, with the RAINBOW orchestrator being part of the Centaurus initiative supported by the Linux Foundation project. Additionally, RAINBOW partners participated in European Big Data Value Forum 2022 while also the consortium conducted synergies with nine EU projects (https://rainbow-h2020.eu/clustering/) in order to disseminate its outcomes and influence further the research and industrial community.

To this end, the RAINBOW project itself helps the European providers to remotely manage their infrastructure that is running, potentially, on hundreds of edge devices.

Copyright © Rainbow Consortium Partners 2020 -2022

4 RAINBOW's Research Outputs

In this section we focus on the results of the project's technical work in a way that sheds light on the impact this work has on the strategic objectives of the call. In order to do so, first we remind certain impact KPI's presented already in the project's DoA [1]; many of those are naturally achieved through the features of the main RAINBOW components.

Table 4-1 Impact KPIs linked to the project's research output

Impact KPIs linked to the project's research output
KPI.1 Decrease in effort and investments for developing and managing the lifecycle of fog computing services
KPI.4 Improved efficiency and performance of fog nodes due to more efficient service development and orchestration
KPI.5 Increased productivity of business applications which rely on, or can be developed based on fog computing services
KPI.6 Increased trust feeling of data and services relying on fog-based services and IoT infrastructure in general and improved security and privacy guarantees
KPI.8 Lowering access barrier for SMEs for to the usage of advanced technologies for the development, management, security and orchestration of services, made available to SMEs

Starting from the updated view of the conceptual architecture depicted in section 2.1.2 of D5.3 [3], we concentrate on those components which directly affect the KPI's above.

First in line are the **client-side tools of the modelling layer** (Service Graph Editor, Policy Editor and Analytics Editor). These tools are accessed through the Dashboard and help achieve KPI.1 and KPI.8. The Dashboard contributes to the aforementioned KPIs on the one hand, by reducing the needed effort and time for the deployment procedure; on the other hand, during the whole lifecycle management of the applications by reducing the complexity of their on-boarding and by providing the graphical tools for their overall management. Additionally, client-side tools contribute to these KPIs by enabling the end users to have access and utilize all the RAINOW features, as much as possible regardless

Copyright © Rainbow Consortium Partners 2020 -2022

to their previous working experience related to that feature, through graphical abstractions and enforcement of a standard procedure whenever they need to utilize them. Additionally, through the RAINBOW Dashboard, end users could visualise various details related to the very deployment, such as the status of each part of it, information that relates to the monitoring and security of the deployment, as well as other running details that can be depicted or managed through a user interface that the user must have access to.

The **RAINBOW Orchestrator** contributes to KPI.1 by reducing the deployment and maintenance effort for fog applications using the Service Graph as a universal application deployment description. Upon submission of a Service Graph by the user, the Orchestrator generates all Kubernetes-native deployment and configuration objects to deploy the application, as described in the high-level Service Graph abstraction. Additionally, RAINBOW-specific objects, i.e., Service Level Objective configurations, are generated. Upon update of an existing Service Graph, the Orchestrator proceeds as necessary to ensure that the system state matches the specification in the Service Graph. This greatly reduces the development and maintenance effort required by developers of fog applications, who no longer need to deal with low-level deployment configurations.

The **RAINBOW Scheduler** contributes to KPI.4 and KPI.5 by ensuring that service instances are placed on the compute nodes that are most suitable for them in terms of hardware requirements and network conditions. The Scheduler is triggered whenever a new service instance needs to be deployed, either due to a Service Graph addition/change or due to scaling. The node chosen for each service instance needs to comply with the service's hardware and network Quality of Service (QoS) requirements. While hardware requirements compliance is also provided by the default Kubernetes scheduler, the network QoS compliance provided by the RAINBOW scheduler greatly increases the productivity of fog applications, by ensuring that they have the necessary network bandwidth, latency, etc. that they require to perform as desired. Additionally, by allowing the collocation of service instances on already used nodes, the number of active fog nodes can be reduced to save battery and costs, thus improving node efficiency.

Monitoring functionalities of the RAINBOW ecosystem increase the productivity of the end-users by alleviating difficulties related to Fog nodes, service, and application-related monitoring. By installing the RAINBOW platform, users can inspect (out-of-the-box) the state of their infrastructure without any extra configuration or installation (KPI.1). Moreover, the RAINBOW monitoring agents minimize monitoring data transfer, size, and energy consumption by providing adaptive interfaces for monitoring and measurement dissemination, contributing to RAINBOW KPI.4. Furthermore, an open-source monitoring library for python programs along with its extensible interfaces allows the user easily to expose application-level metrics. The latter is crucial for SLOs' definition, and other monitoring systems either do not provide it at all or require programming effort from the developers of the fog applications (contributing to KPI.5).

The **Data Management & Analytics** component enables the storage of the monitoring data, performs on-time streaming analysis on them, and facilitates the SLOs' definition.

Copyright © Rainbow Consortium Partners 2020 -2022

Thanks to RAINBOW, users can submit their streaming queries through the RAINBOW dashboard in a declarative way (KPI.1) and the RAINBOW platform translates the queries into efficient streaming distributed execution minimizing specific objectives, like energy consumption, latency, accuracy, etc. (KPI.4). Through that, RAINBOW increases the observability with users having a better understanding of the underlying Fog infrastructure and running services (KPI.1, KPI.5). More importantly, the SLO controller utilizes the analytics stack to perform monitoring analysis on time and consequently alter the deployment accordingly (KPI.1).

The **Service Level Objectives (SLOs)** and **Elasticity Strategies** contribute to KPI.5 by ensuring that applications operate within the desired performance bounds without incurring excessive costs. In addition to fluctuating user requests, as known from cloud environments, the fog present additional challenges for applications, such as varying network performance and performance requirements that derive from IoT-based scenarios. The flexibility of the SLO controllers provided by RAINBOW allows capturing simple and complex SLOs, which are then automatically enforced at runtime through Elasticity Strategies to ensure the applications' performance goals are met.

The **RAINBOW attestation toolkit** comprises three opt-in security enablers to fog application developers with privacy preservation and remote orchestration in mind, thus helping to run and orchestrate services that mix in multi-domain environments within the RAINBOW ecosystem, strongly contributing towards KPI.6.

With Configuration Integrity Verification (CIV), we utilize sound, trusted computing technologies (e.g., the prominent Trusted Platform Module) as trust anchors in edge devices that subsequently allow edge devices to report verifiable evidence about the correctness of the platform state. Users of the platform can use this functionality to ensure that edge devices are securely onboarded (based on evidence about the device's platform configuration state) before being permitted access to the RAINBOW services. Furthermore, by having the CIV security enabler running continuously, the devices can continue to give assurance about their configuration correctness throughout their entire lifetime.

With privacy preservation in mind, the novel protocol underlying CIV requires no concrete platform information ever to leave an edge device when asked to prove its correctness, thus allowing users that orchestrate even the most privacy-critical services to benefit from the CIV functionality's security guarantees without having to sacrifice privacy in the process. Furthermore, the CIV protocol can be complemented with the Direct Anonymous Attestation (DAA) component (also part of the open-source attestation toolkit) to make any digital signatures the device presents unlinkable.

For users who want evidence about the correct execution of highly security-critical services running on edge devices, the Control-Flow Attestation (CFA) functionality gives even the most resource-constrained devices that execute deeply embedded programs the ability to prove that programs (or segments thereof) executed correctly on the device (i.e., that a program completed without errors and in the absence of any runtime attacks).

Copyright © Rainbow Consortium Partners 2020 -2022

The novel CFA protocol developed for RAINBOW leverages state-of-the-art zkSNARK-based Verifiable Computation techniques to deliver such program execution correctness guarantees in a cryptographically sound manner.

4.1 Impact on Demonstrators

This section is dedicated to highlighting how the provided RAINBOW functionalities were able to help add value to the operations of the three advanced fog use cases, i.e. the three demonstrators described in the project. In addition, we point out the feedback directed from the demonstrators back to the technical teams, which sketches the advantages and the disadvantages the developed solution has.

The target is to promote the adaptability of the RAINBOW platform on a European level, effectively supporting the modernisation and the increased efficiency of the involved industrial stakeholders' work, considering the corresponding stakeholder requirements, as identified, and documented in D1.1 [9].

4.1.1 Human-Robot Collaboration in Industrial Ecosystems

While using the early release of the RAINBOW platform, the first demonstrator noticed that the RAINBOW UI significantly simplifies the deployment process, as the end-user does not have to actually write, configure and re-edit YAML files for successful deployment. Furthermore, the early release of the RAINBOW platform provided a first version of service graph visualisations which indicate interdependencies, deemed to be very useful to visualise complex application deployments.

At the same time, the early release of the RAINBOW platform was found to be lacking automated deployment scripts to achieve the deployment of the RAINBOW platform at the demonstrator's infrastructure in a zero-touch manner. Therefore, platform end-users had either to rely on pre-existing highly technical knowledge or needed to be assisted by a technical counterpart during the deployment process. Additionally, another reported issue was the lack of documentation which would help understand how the deployment of the components would look like on a container level; this made it difficult to notice issues and debug them appropriately. Furthermore, the deployment of the early release of the RAINBOW platform identified that there existed no support for volume mounts/config maps, making it hard to upload application-specific configurations. Finally, it was noted that the RAINBOW UI does not provide sufficiently detailed logs and does not allow direct SLO configuration updates. Lack of sufficiently detailed logs causes issues related to debugging, while lack of direct SLO configurability means that the end-user has to engage in such modifications manually using the command-line interface (CLI) in the respective virtual machines on which the services at hand are running.

When deploying and testing the advanced release of the RAINBOW platform, the demonstrator added to the advantages of the early release, first the enhancement of most of the positive points highlighted during the early stage and in particular the service graph visualisation. Furthermore, according to the first demonstrator the deployment of

Copyright © Rainbow Consortium Partners 2020 -2022

the RAINBOW platform has become much simpler with the help of convenient and easy-to-use scripts, reducing the efforts and time significantly. This is aided by the newly created documentation on the installation and deployment process, which was deemed very helpful and enabled a zero-overhead process. Moreover, using RAINBOW increases the ease and speed of deployment of use-case specific applications through an intuitive web UI. Continuing, RAINBOW now provides secure communication between the components by integration of CJDNS to RAINBOW ecosystem. Lastly, the RAINBOW platform now supports the creation of SLOs for out-of-the-box metrics (for example *cpu_ptc*) through the web UI, even though custom metrics are yet not supported.

As for the weak points identified in the advanced release, the first demonstrator identified the weakness of missing support for creation of SLOs using the custom metrics of the applications through the web UI. This constitutes a weakness of the RAINBOW platform as it necessitates more technical knowledge, as well as the usage of YAML files to instantiate and configure certain SLOs. Continuing, in particular for the first demonstrator, attempting to deploy an application out of order or in a random fashion, may result in errors and subsequently lead to unsuccessful deployment. Assuming the deployment requirements of the first RAINBOW use case, the rabbitmq and redis services should deploy before other pods.

Constructive feedback for further improvement of the RAINBOW platform has been identified by the first demonstrator at the end of their involvement in the project. RAINBOW offers no support for using of volumes for the containers. This significantly limits the interaction and configuration opportunities with the deployed pods, which is highly used for containerized workloads. Many RAINBOW artifacts are deployed using "Network Mode Host" option; this option has proved to enable more access than necessary or desired to the overall server, which may potentially raise security concerns.

When it comes to the Human-Robot Collaboration use case, the most impressive improvements are:

- Robots successfully stopped when latency exceeds 500ms - adjustable through OPCUA SLO controller (please refer to BIBA-KPI-01 of D6.9 [16]).
- RAINBOW provides secure mesh overlay network out-of-the-box; no security incidents whatsoever were reported (please refer to BIBA-KPI-11 of D6.9 [16]).

4.1.2 Digital Transformation of Urban Mobility

While using the early release of the RAINBOW platform, the second demonstrator noticed that with the help of RAINBOW UI, the process of component and application creation simplifies the deployment when comparing it with a pure Kubernetes deployment, thanks to a higher-level approach that does not require any knowledge about Kubernetes

Copyright © Rainbow Consortium Partners 2020 -2022

and its technical files for successful deployment. Moreover, the RAINBOW platform provides a service graph which is a very useful method to visualise service dependencies and the whole deployment architecture.

At the same time, the early release of the RAINBOW platform was found to be lacking a fully automated installation script, similarly with the first demonstrator. An issue particular to the second RAINBOW demonstrator was the additional complexity introduced by the need to be able to migrate a set of services between nodes with ARM64 and AMD64 processors.

When deploying and testing the advanced release of the RAINBOW platform, the demonstrator added to the advantages of the early release, the SLO framework of RAINBOW which allows specifying complex SLOs that can be used in conjunction with multiple elasticity strategies to adapt deployments at runtime. Furthermore, the monitoring framework of RAINBOW was pointed out as yet another strength of the platform; it was observed that implementing the same logic for a migration with application-level metrics without RAINBOW would have been difficult or even completely impossible. Lastly, arguably the most important demonstrator-specific improvement was the capability to migrate services between nodes with ARM64 and AMD64 processors.

As for the weak points identified in the advanced release, it was reported that error messages appear to be somewhat generic and don't offer much insight to the underlying issue. It is thus important to improve the error message management feature, as the current one is too generic and often it has been necessary for the second demonstrator to contact the technical counterparts to resolve potential issues. Lastly, the documentation can be improved to support the development and debugging process, in order to avoid consuming additional resources in remote assistance.

Copyright © Rainbow Consortium Partners 2020 -2022

Constructive feedback for further improvement of the RAINBOW platform has been identified by the second demonstrator at the end of their involvement in the project. At first, this demonstrator considers that external volumes binding can drastically speed up and improve the RAINBOW experience since it would allow dynamic configuration parameter changes on the fly, while the host network mode should not be mandatory, since it may create security issues in production deployments. Also, the user interface could be extended to enable the end user to effortlessly edit the SLOs.

When it comes to the Digital Transformation of Urban Mobility use case, the most impressive improvements are:

- Average fog node execution power consumption was approximately 30W; edge node execution average power consumption was approximately 5W; average power consumption was less than 10W, which is not only in line with the expectations but significantly exceeds them (please refer to AHED-KPI-01 of D6.9 [16]).
- The expected time for a successful deployment was less than 1 min; RAINBOW achieved a deployment time of 20 sec in its first release, and a deployment time of 10 sec in its final release (please refer to RAINBOW-KPI-01 of D6.9 [16]).
- The software delivery cycle when using RAINBOW's advanced release was about 10 seconds, whereas the expected value was 1 minute (please refer to RAINBOW-KPI-02 of D6.9 [16]).

4.1.3 Power Line Surveillance via Swarm of Drones

While using the early release of the RAINBOW platform, the third demonstrator identified several strong points regarding the RAINBOW platform. Tests confirmed that RAINBOW, by enabling the direct use of distributed services, allows for more efficient operations without the efforts of developing and implementing them in order to change the existing centralized services. Furthermore, while testing the early release it was noted that the usage of SLOs allows for a simpler physical topology of the system. Thanks to the usage of SLOs, there is no need to have different types of nodes (i.e., all hardware can be uniform), so operators can focus on the application layer of their tasks without being concerned with the condition of the system and proper placement of hardware. The orchestrator automatically assigns services to proper nodes and makes sure that this allocation is optimal while nodes are added and removed from the system; this is pivotal to the demonstrator and showcases directly applicable added value for the demonstrator.

When deploying and testing the advanced release of the RAINBOW platform, the demonstrator observed the enhancement of most of the positive points highlighted during the early stage and in particular drone operation efficiency, thanks to the successful coordination of the ground control stations enabled by RAINBOW. Tests conducted by the third demonstrator revealed a noticeable increase in the efficiency of

Copyright © Rainbow Consortium Partners 2020 -2022

the drone system missions; this is attributed to the good coordination of the work of the GCS station, enabled by the underlying RAINBOW components. After the enhancements and improvements undergone, the installation and instantiation of the use case scenarios became more straightforward. For all intents and purposes, all deployment-related activities have been completely automated by the provided scripts. It is also important to note that SLOs can now be used to adjust system performance and to make system deployment adapt automatically to the available hardware. This feature is of utmost importance for the third demonstrator, as its success is directly dependent on the flawless cooperation between networked elements of varying hardware typologies and resource availability.

Constructive feedback for further improvement of the RAINBOW platform has been identified by the third demonstrator at the end of their involvement in the project. It is advisable to develop a tool that will allow users to assess the state of RAINBOW cluster, to check whether installation has succeeded and whether all components are working as expected with just a single command. Additionally, the third demonstrator found out that installation of the analytic stack requires tweaking of configuration files and the documentation does not state clearly whether IPv6 or IPv4 addresses should be used; this is not as much a functional issue as much a documentation shortcoming. Continuing, overhead if the installation script for the master node would output configuration parameters in a form that can be directly pasted into installation script for worker nodes (or as a configuration file that this script reads), it would significantly reduce time. Another hint provided by the third demonstrator is that a declarative approach to installation should be considered (for example Ansible), to allow for an easier and more application-focused deployment process. Furthermore, although RAINBOW now supports custom metrics, they cannot yet be treated the same way as those built-in; this could be addressed to allow for a more application-aware service instantiation. Lastly, when a horizontal SLO scales up a service deployment, the client services still use only the previously running instances; when that is the case, it would be very beneficial for the users if they could detect in some easy way that there is a new instance available and they should reconnect, so that the load is spread also across the new instance.

When it comes to the Power Line Surveillance via Swarm of Drones use case, the most impressive improvements are:

- The measured increase of productive flight distance per drone was 100% in the second release, with the initially expected value being 50% (please refer to MSP-KPI-02 of D6.9 [16]).
- Reduction of data acquisition time per kilometer of power line was 61.2% in the second release, exceeding the expectations of a 50% increase (please refer to MSP-KPI-03 of D6.9 [16]).
- RAINBOW's advanced release managed to increase the efficiency of battery usage for a productive phase of the drone flight by 167.2%, significantly exceeding the expected value of 55% (please refer to MSP-KPI-05 of D6.9 [16]).

Copyright © Rainbow Consortium Partners 2020 -2022

4.2 Lessons Learnt

Successes and failures resulting from testing the early release of the RAINBOW platform at the premises of the three demonstrators were turned into lessons learnt, improvement points, as well as valuable data concerning the overall experience of the demonstrators and resulted into an improved advanced release of the RAINBOW platform.

While the early release of RAINBOW was tested using virtual deployments of the three demonstrators (e.g. lab-to-lab set-up), the advanced release of the RAINBOW platform was deployed and tested on physical deployments of the three demonstrators [12] [13] [14]. The three use cases have been set-up in operational environments, the RAINBOW platform and use case specific applications have been deployed on physical edges and the testing and collection of metrics was performed while the physical demo was operational [16]. Given the definitions of Technology readiness levels (TRLs) as a method for estimating the maturity of technologies, **RAINBOW features a TRL of 7 "System prototype demonstration in operational environment"**. Regarding the modules developed by the project partners, the TRL level achieved by each module at the end of RAINBOW project along with the owner of the respective IPR, a separate IPR inventory is maintained and presented in D7.9 [11].

In this subsection, matters concerning implementation, operation and execution of the demonstrators are formulated as methodological lessons learnt for the further exploitation, effective utilisation, and further development of the RAINBOW platform.

An important request from the demonstrators from the early release was the automation of the deployment processes and the disengagement of human operators from virtual infrastructure management. This was envisaged to better align RAINBOW with the stakeholders' expectations. Enabling this functionality proved to indeed be a pivotal enhancement, and raised yet another desire, namely the configurability of the automated deployment via dedicated config files.

What end users value in RAINBOW:

1. Automation of containerised application deployments is a strong point, highlighted by all three demonstrators. Throughout both development stages, end users noted the added value offered by the fact that RAINBOW made it measurably easier to deploy highly complex services without having to deal with low-level configurations. This eliminates knowledge barriers and enables relevant stakeholders to act more independently. Especially during the final stage, RAINBOW allowed for a completely automated installation process using convenient scripts, which boosted this observed advantage.
2. RAINBOW effectively creates an abstraction layer between the infrastructure and the end-user through the RAINBOW dashboard. The provided dashboard significantly simplifies service deployment and component creation. RAINBOW manages to conceal the Kubernetes details and the service graph

Copyright © Rainbow Consortium Partners 2020 -2022

which is appreciated by the end users for managing the dependencies and the whole deployment architecture. Further pursuing the separation of the control and user planes seems to be highly desired for enabling wide market adoption.

3. Automatically adjusting performance and adapting to available hardware proved to be highly desired. Highlighting the newly introduced configurability and overall elasticity of RAINBOW is the fact that it now supports migration from and to nodes with different processor architectures. As such, this adoption guideline incorporates configurability and hardware-agnosticism to steer RAINBOW into a more adaptable direction, due to its capacity to remove overhead from end-users and support a more business-centric approach.

RAINBOW has been built by the technical partners towards the direction set by the elicited stakeholder needs, namely: being **business-oriented** and **user-friendly**, offering **simplified deployment of complex applications** and easy-to-understand **visualisations of interdependencies and application architectures**.

Building on the received feedback and to push adoptability even further, it has been successfully pivoted by the technical partners towards enabling **complete deployment automation** whilst ensuring **easy configurability**, supporting the creation of an **abstraction layer between user and infrastructure**, ensuring **elasticity and adaptability** in terms of hardware and node typologies.

What end users would like to see from RAINBOW:

1. Handling and managing error messages need improvement, along with the corresponding documentation which appears to be lacking in terms of supporting the debugging process and clarifying expected user input. Tests by the industrial partners seemed to hint that a more instruction-based approach could reduce the need for technical partners to intervene in the case of debugging-related issues.
2. End users value the possibility to easily assess the state of the RAINBOW cluster being used. As a guideline for additional developments, enabling end users to easily and timely receive feedback on the status of undergoing deployment processes.
3. Stakeholders appreciate the possibility to use custom metrics for the creation of SLOs but would like to see such metrics receive the same "treatment" as those inherently contained within RAINBOW. It would lead to wider adoption of the RAINBOW platform if it were possible to use the web user interface to create SLOs from custom metrics and not necessarily rely on modifying configuration files.

Considering what RAINBOW has already achieved in terms of configurability and ease of use, the above-mentioned remarks serve as a future roadmap. Overall, RAINBOW will

Copyright © Rainbow Consortium Partners 2020 -2022

ensure even further market penetration and industrial adoptability through **support for user-defined metrics for SLOs, monitoring capabilities,** and **clear documentation.**

Copyright © Rainbow Consortium Partners 2020 -2022

5 Research, Policy, and Industry Community Benefits

5.1 Contribution to Standardisation

According to the DoA, OASIS TOSCA[2] was supposed to be targeted. It is an open standard for modelling the topology of cloud applications, including the relationships between their components and the management processes involved. The original plan for RAINBOW was to extend TOSCA for the fog, thus making an official contribution to this standard a logical step. However, we decided to take a Kubernetes-native approach for the design and development of RAINBOW. Kubernetes is one of the most popular container orchestration systems in production use, which means that support for this platform leads to higher chances of adoption of RAINBOW by the open source and industrial community.

Since, RAINBOW does not make use of TOSCA, it no longer makes sense to contribute to this standard, leading to the exclusion of TOSCA from our list of target standardization groups.

5.1.1 Contribution to TCG

The Trusted Computing Group (TCG)[3] is a not-for-profit organization formed to develop, define, and promote open, vendor-neutral, global industry specifications and standards, supportive of a hardware-based root of trust, for interoperable trusted computing platforms. TCG's core technologies include specifications and standards for the Trusted Platform Module (TPM), Trusted Network Communications (TNC) network security and self-encrypting drives.

IFAT participates in TCG amongst others as President and Chair of the Board of Directors, as Chair of the Certification Program Committee (CPC), which targets especially the development of the new TCG Compliance and Conformance program, and the TSS working group and as Co-chair of the Security Evaluation working group.

RAINBOW contributed to the TCG standardization with related organizational and technical support, technology, and know-how, and discussed the latest advances in DAA and attestation/revocation that have been performed by RAINBOW.

5.1.2 Contribution to Fog Adoption IFG

RAINBOW established the Fog Adoption Industrial Focus Group (FA-IFG) which counts currently eight members. The FA-IFG shall mainly act as information hub for its members coming e.g. from Digital SMEs and Start-ups, Cloud and IoT services, and platforms providers. The target audience are Fog-enabled application developers and end-user

[2] https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=tosca
[3] http://www.trustedcomputinggroup.org

Copyright © Rainbow Consortium Partners 2020 -2022

organizations, as well as Fog and Cloud experts, companies, and institutes, but also researchers, experts, scientists and individuals interested in Cloud computing. On the RAINBOW website, a dedicated section about the FA-IFG has been implemented where the members are listed with their picture and some more information. In the course of the RAINBOW project, the FA-IFG members have been informed and involved by e.g. personal invitations to webinars and workshops offered by RAINBOW. The post-project activities of the FA-IFG will stay in responsibility of the Sustainability Board defined in the Grant Agreement.

5.1.3 Contribution to the Open Horizon initiative

We have presented RAINBOW to Open Horizon's Technical Steering Committee and discussed the possibility of contributing RAINBOW components to Open Horizon. In the end, the Fogify emulator and the Polaris SLO Framework were selected by the RAINBOW consortium and offered to Open Horizon. Details of this activity are presented in deliverable D7.5 [8].

5.1.4 Contribution to the Linux Foundation's Centaurus project

We have managed to contribute a fork[4] of the RAINBOW scheduler to the Linux Foundation's Centaurus project[5]. Centaurus is a novel open-source platform for building unified and highly scalable public or private distributed cloud and edge systems. Details of this activity are presented in deliverable D7.5 [8].

5.2 Alignment with OpenFog architecture

RAINBOW conforms to and builds upon the "Economic" pillar of the Lisbon strategy as it accelerates the adoption of IT applications and services in domains where operations must be performed at the edge nodes, improving the competitiveness of the European software and industrial sector and promoting growth.

As per the DoA, the project aspired to liaise and contribute to the results of the **OpenFog Consortium**. The OpenFog was a consortium of members extending from academia to industry innovators (i.e., Princeton, ARM, Cisco, Microsoft) that signed a memorandum of understanding in 2016 with the goal of fostering interoperability across the fog continuum to drive IoT. During the project, in 2019, the OpenFog consortium merged with the Industrial Internet Consortium[6]. Before the merger, the OpenFog consortium released in 2017 a seminal white paper entitled "OpenFog Reference Architecture" that outlines the pillars of an open, modular and fog computing ecosystem.

[4] https://github.com/CentaurusInfra/polaris-scheduler
[5] https://www.centaurustcloud.io
[6] Rebranded as the **Industry IoT Consortium** in August 2021.

Copyright © Rainbow Consortium Partners 2020 -2022

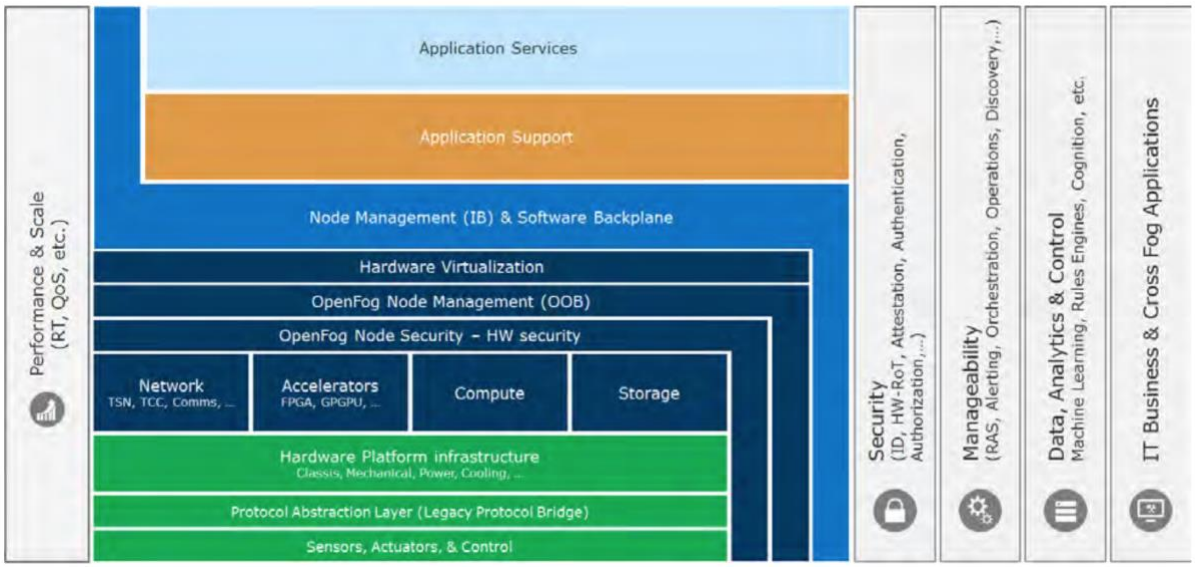

Figure 5-1 OpenFog Reference Architecture and Perspectives

Figure 5-1 presents the OpenFog reference architecture (as introduced in the respective white paper). In this figure, the fog continuum with all that such an ecosystem comprises (i.e., hardware, software, applications) is depicted horizontally, while 5 perspectives extend vertically and denote what fog service platforms can bring to the table when offering services to users moving to the fog continuum. It must be noted that RAINBOW adopts the Open Fog architecture template, not altering the horizontal fog continuum, and offering its services as vertical enrichments to the continuum. No constraining requirements to the fog continuum constituted a particularly important design decision for RAINBOW as this would have implications to the suitability of applications that can embrace the RAINBOW ecosystem and narrow the reach and target audience of the platform in an area where heterogeneity across the continuum is the norm and not the exception. Hence, for RAINBOW, users may bring-their-own edge devices, fog nodes and public cloud offerings, while from an application perspective no alterations to the business logic are required; RAINBOW limits its requirements so that it is enough for IoT applications to adopt the microservices architectural paradigm with application services containerized by exploiting either Docker or the OCI specification.

In terms of vertical offerings to the fog continuum, RAINBOW makes contributions to three of the OpenFog perspectives, namely to (i) Security; (ii) Manageability; and (iii) Data Analytics and Control. In turn, RAINBOW extends the number of perspectives in the architecture specification to include a 6th perspective in an area of importance, especially, for the next generation of data intensive IoT services. The 6th pillar is named rapid prototyping and testing, and RAINBOW's contribution to this perspective includes the Fogify emulator for geo-distributed IoT services. The following figure presents the RAINBOW services in accordance with the OpenFog architecture specification.

Copyright © Rainbow Consortium Partners 2020 -2022

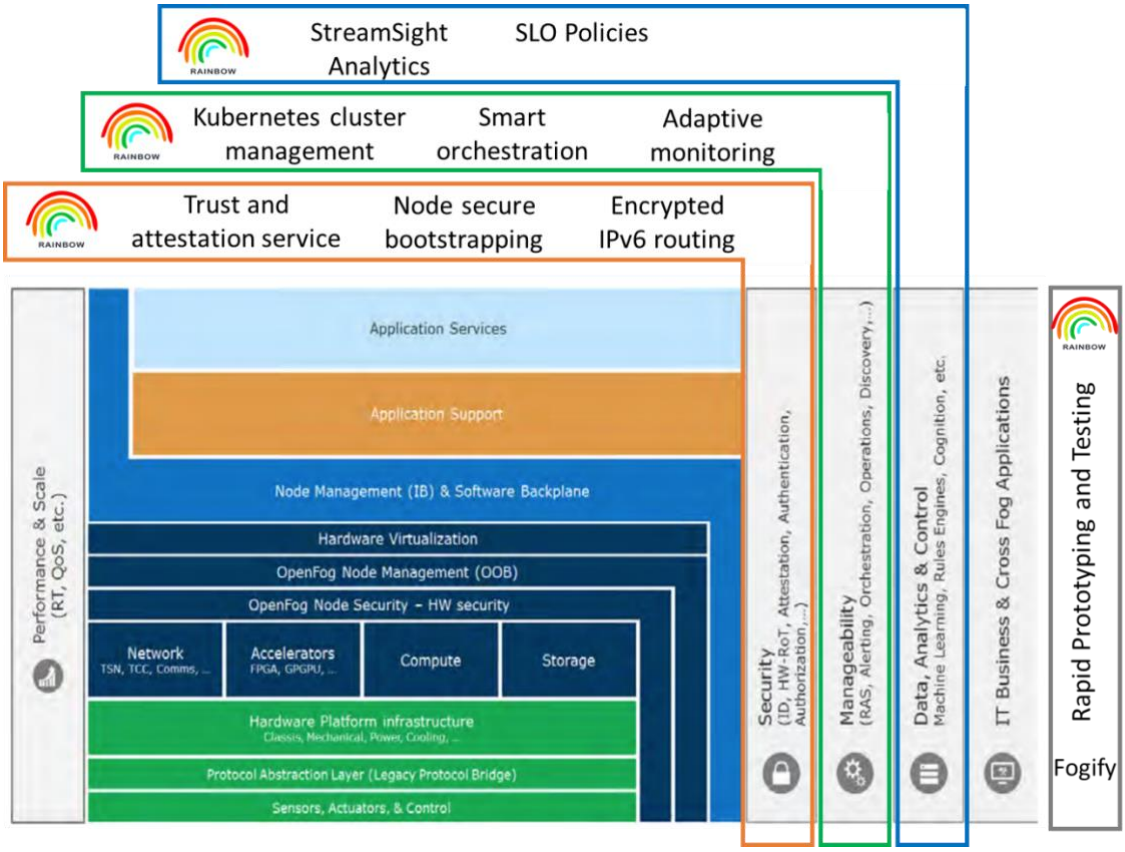

Figure 5-2 RAINBOW and OpenFog superposition

With relevance to security, RAINBOW contributes to the OpenFog security pillar for device and software integrity, trust and attestation, measurement, and verification. Specifically, the RAINBOW trust and attestation service handles the secure onboarding of new nodes to the cluster by assuring their trustworthiness through attestation practices on their OS and RAINBOW related add-ons; the vulnerability assessment mechanism identifies potential vulnerabilities of the deployed containers and notifies the end user in case something found. The node secure bootstrapping works together with the trust and attestation service to verify the correctness of the RAINBOW plugins and assure that the bootstrap of RAINBOW is correct without any security incidents. Finally, the encrypted IPv6 networking aids secure connection and data transmission between the nodes, by utilizing some cryptographic algorithms and cryptographic keys that are generated through the corresponding RAINBOW service.

With relevance to manageability, RAINBOW provides services that address the scalability, autonomy and reliability pillars by offering through its dashboard and service graph editor the ability to design, configure, monitor and govern large-scale geo-distributed deployments. These features cover Kubernetes cluster management, software defined networking, monitoring the resource utilization of underlying fog nodes, network overhead and enabling the deployment of custom app-level monitoring probes, and triggering vertical/horizontal scaling actions. When combining the aforementioned services with the RAINBOW analytics services, the fog servicing

Copyright © Rainbow Consortium Partners 2020 -2022

ecosystem can cover the OpenFog autonomy and agility pillars as users have the ability to define and extract high-level analytic insights from raw monitoring data and then add, even at runtime, SLOs which automatically trigger scaling actions to ensure the reliability and performance of their applications.

5.3 Improving Innovation Capacity

Towards this goal, the RAINBOW consortium has submitted a project proposal to the **Eclipse IoT community**. The setup of Mobile Ad-Hoc Networks is currently not sufficiently addressed by the Kubernetes community, yet these networks are very common in fog computing. Thus, RAINBOW has submitted an Eclipse IoT project proposal for a Secure Admission Control Protocol for Kubernetes in Mobile Ad-Hoc Networks. Details on this activity are presented in deliverable D7.5 [8].

5.4 Contribution to European Policies and Other Strategic Initiatives

RAINBOW contributes to all three EU **Digital Single Market (DSM)** pillars, namely to "Access" pillar by lowering the barrier for the development and deployment of advanced cloud services, to "Environment" pillar by creating an environment of trust in the cloud for SMEs and Public Sector Organisations and finally to the "Economy & Society" pillar though offering a solution that will improve interoperability, will contribute to standards, and allow ICT SMEs to grow.

The consortium submitted the **Horizon Results Platform (HRP)**, two exploitable assets of RAINBOW, namely the RAINBOW platform and Fogify (see section 6.2.2), to disseminate the work of the project and foster synergies with other projects, partners or interested investors. The two entries in the HRP were based on the exploitation analysis of these two Key Exploitation Results that the RAINBOW Project delivered. Some excerpts from their submission in the HRP are demonstrated below.

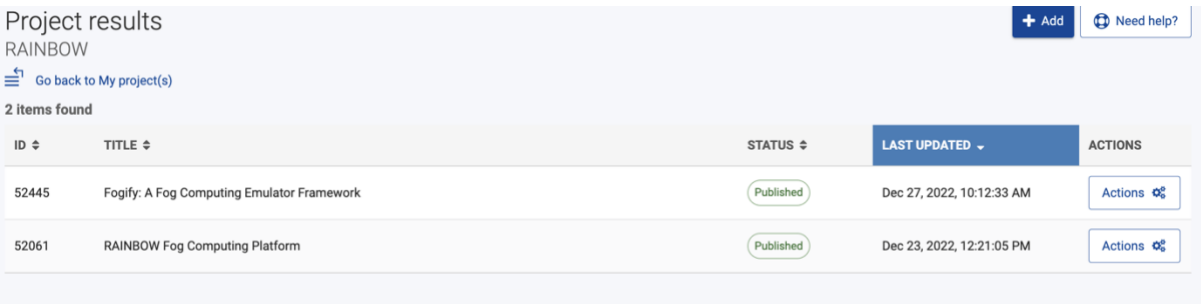

Figure 5-3 Horizon Results Platform – RAINBOW results

Copyright © Rainbow Consortium Partners 2020 -2022

5.4.1 RAINBOW Fog Computing Platform - HRP Entry

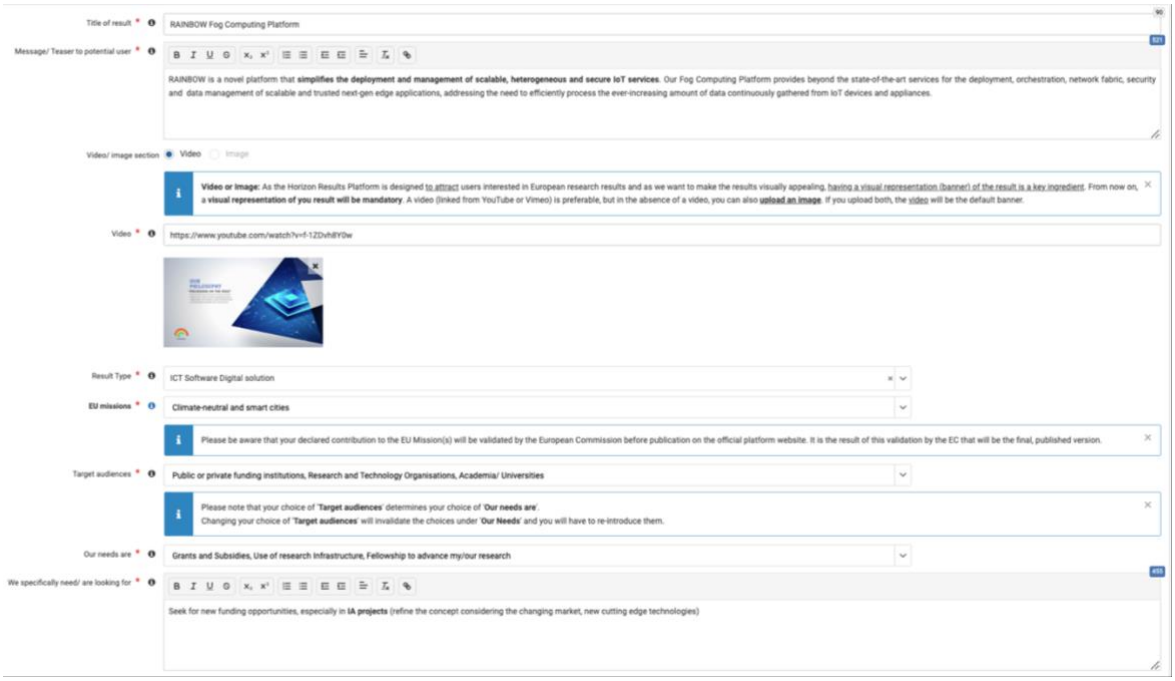

Figure 5-4 HRP - RAINBOW Fog Computing Platform-Result Title, Target Audiences and Needs

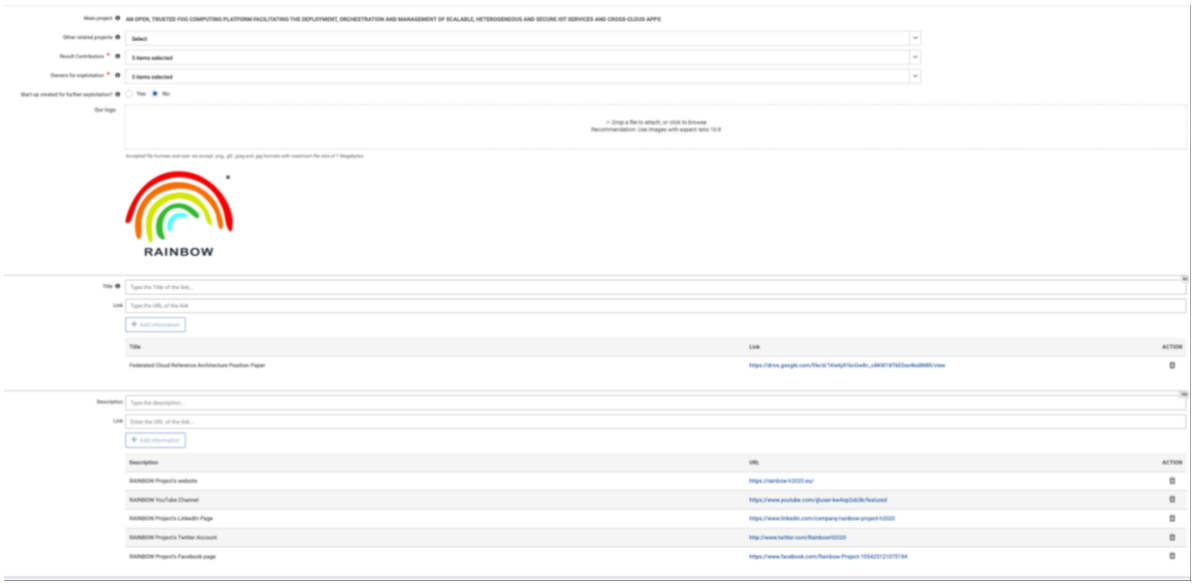

Figure 5-5 HRP - RAINBOW Fog Computing Platform-About us

Copyright © Rainbow Consortium Partners 2020 -2022

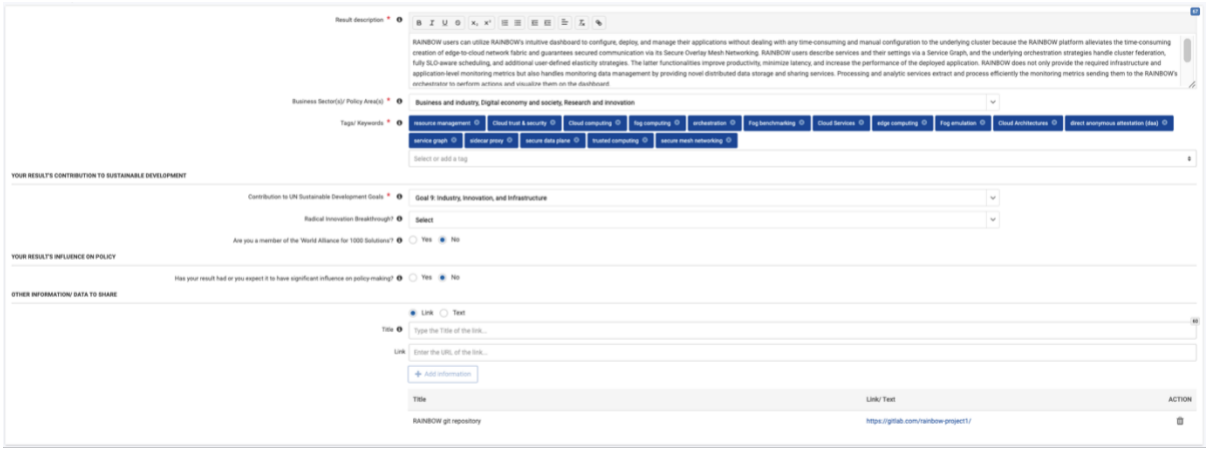

Figure 5-6 HRP - RAINBOW Fog Computing Platform-Result Description and Influence

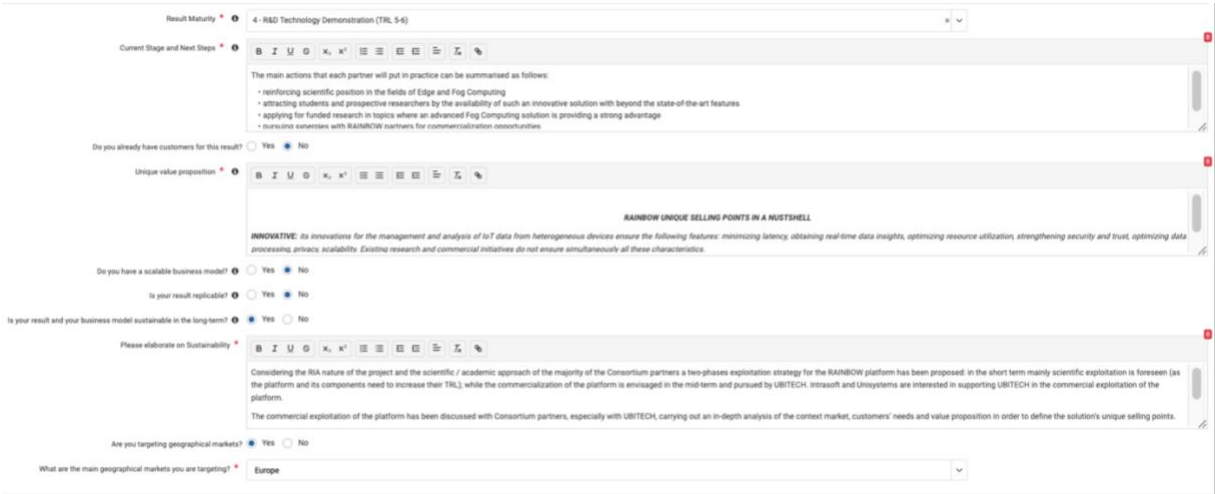

Figure 5-7 HRP - RAINBOW Fog Computing Platform-Result and Business Maturity and Exploitation outlook

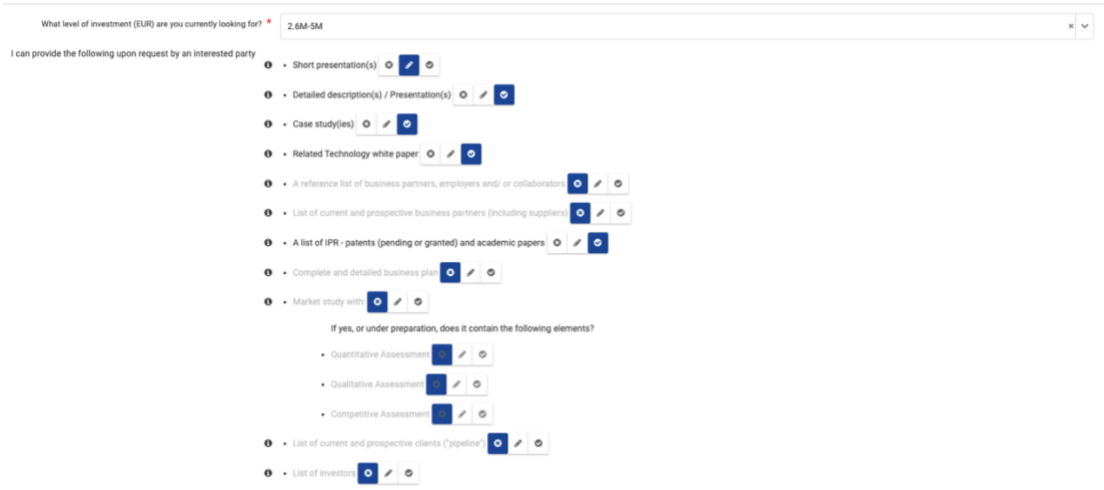

Figure 5-8 HRP - RAINBOW Fog Computing Platform-Investors' Corner

Copyright © Rainbow Consortium Partners 2020 -2022

5.4.2 Fogify: A Fog Computing Emulator Framework – HRP Entry

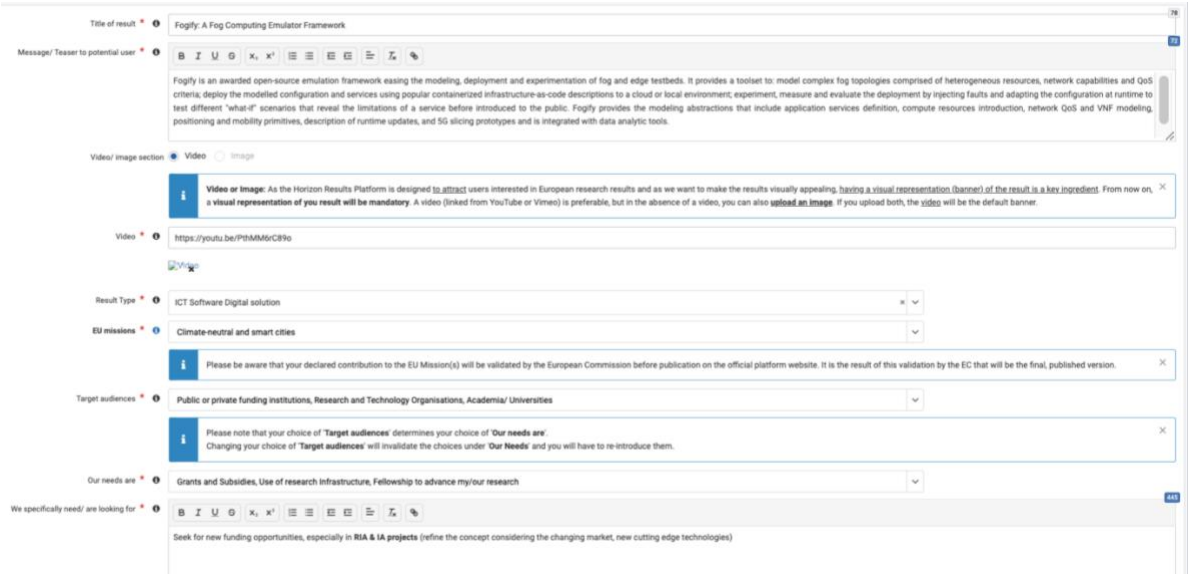

Figure 5-9 HRP - Fogify: A Fog Computing Emulator Framework - Result Title, Target Audiences and Needs

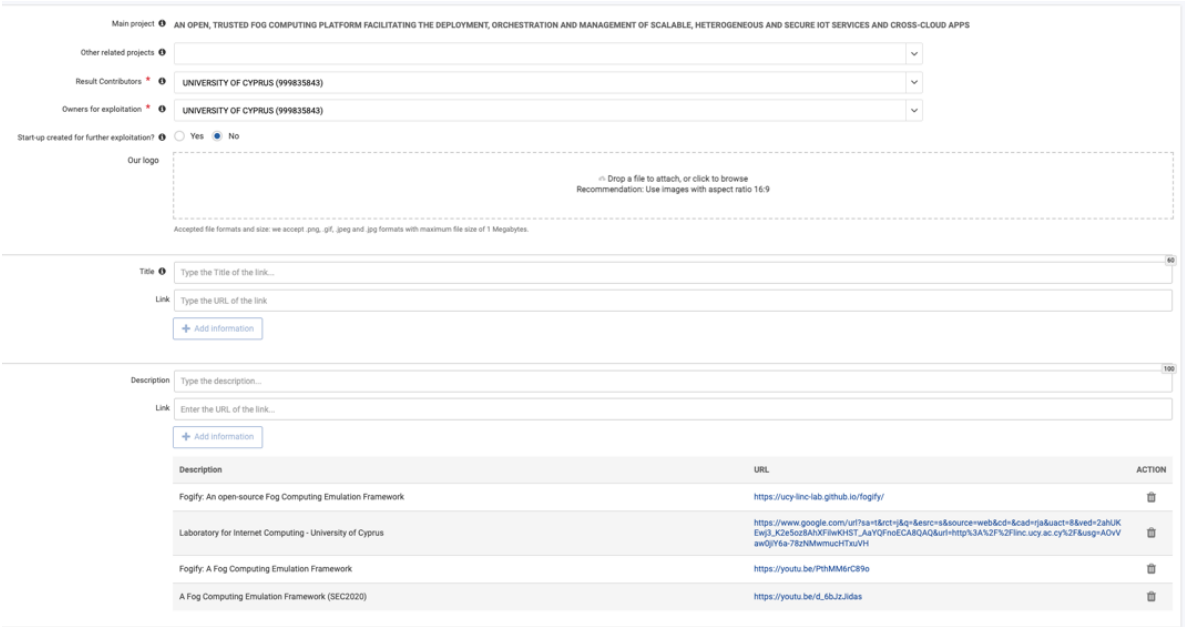

Figure 5-10 HRP - Fogify: A Fog Computing Emulator Framework – About Us

Copyright © Rainbow Consortium Partners 2020 -2022

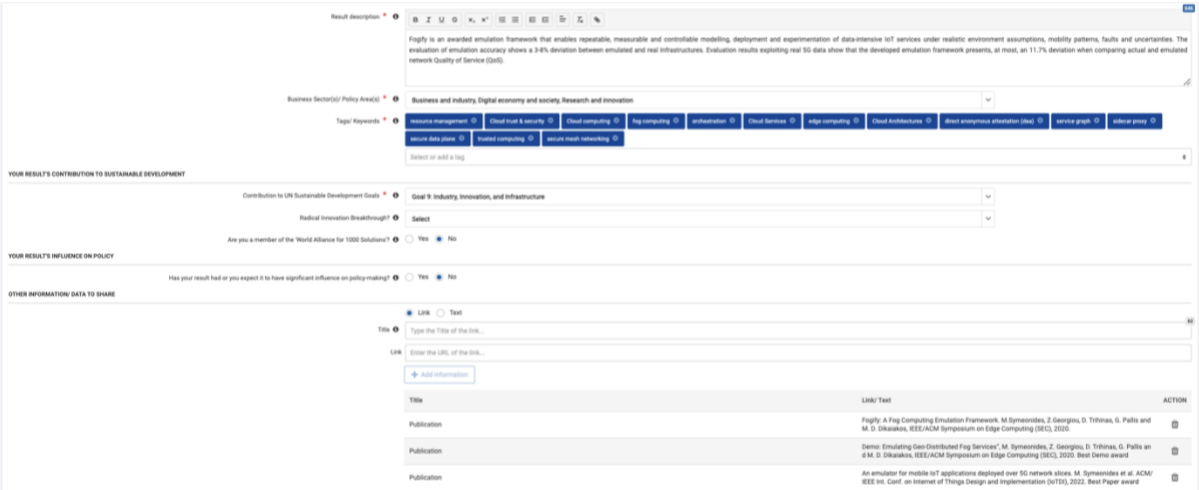

Figure 5-11 HRP - Fogify: A Fog Computing Emulator Framework – Result Description and Influence

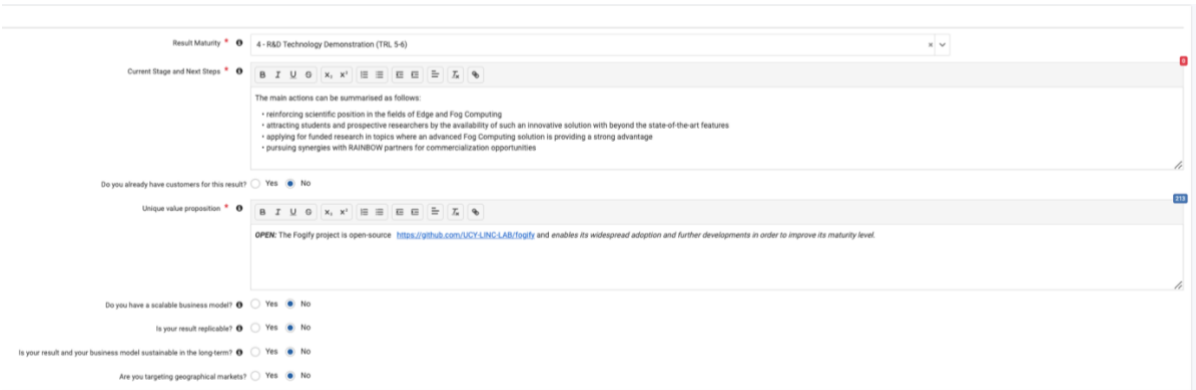

Figure 5-12 HRP - Fogify: A Fog Computing Emulator Framework - Result and Business Maturity and Exploitation outlook

Copyright © Rainbow Consortium Partners 2020 -2022

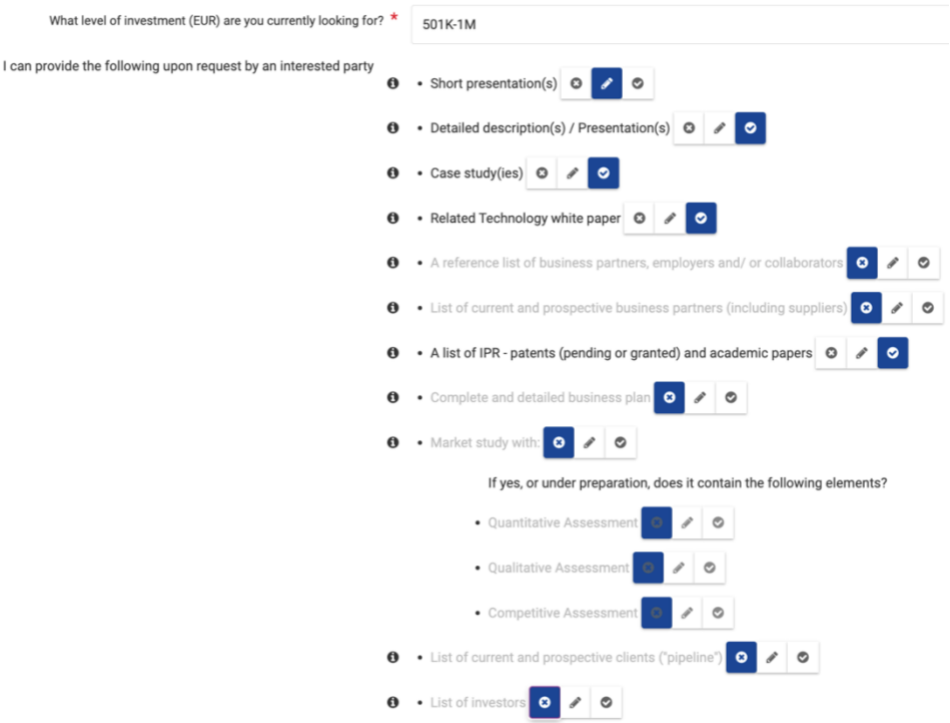

Figure 5-13 HRP - Fogify: A Fog Computing Emulator Framework – Investors Corner

Copyright © Rainbow Consortium Partners 2020 -2022

6 Knowledge Transfer

6.1 Contribution to TPM

RAINBOW has contributed knowledge to the Trusted Computing and Trusted Platform Module ecosystem along several lines.

RAINBOW's formal activities in the Trusted Computing Group, the premier standardization organization for TPM and related technologies is discussed in section 5.1.1 and those towards wider adoption of the Fog paradigm in section 5.1.2.
Applicability and best practises for TPMs and advanced attestation protocols have been presented in a dedicated webinar "Rainbow Security and Trust in For Computing" held by RAINBOW partners IFAT, POLITO and DTU.

RAINBOW also contributed to relevant open-source community, for instance project tpm2-tools on GitHub which provides a practical command interface to the TPM.
RAINBOW is also supporting academic careers around challenges leveraging technologies like the Trusted Platform Module in several PhD theses which are undertaken at partners DTU, POLITO, and IFAT (together with Graz University of Technology, Austria). See also contribution to **ICT-15 Strategic Impact #4** in section 3.

6.2 Dissemination, Communication, Exploitation

RAINBOW has organized a series of impactful dissemination activities which aimed at demonstrating the project outcomes and were focused on the engagement and interaction with targeted groups. Together with the communication strategy, which was based on strong web & social media presence, these activities ensured stakeholders' engagement and the creation of interest in scientific, research and industrial communities.

Members of the consortium organized (or co-organized) a total of five scientific workshops with the support of RAINBOW. All these workshops were focused on scientific topics highly relevant to Edge, Fog and Cloud Computing research and technologies that are part of the core features of the RAINBOW solution.

- RAINBOW with the collaboration of ASSURED, PUZZLE and C4IIoT EU-funded projects co-organized the **3rd Workshop on Cyber-Security Arms Race** (CYSARM) under the prestigious ACM Conference on Computer and Communications Security (CCS, South Korea), on November 19, 2021.
- RAINBOW was part of the joint initiative of EU-funded projects ASTRID, CYBER-TRUST, GUARD, SIMARGL, DATAVAULTS, PALANTIR and SPEAR, that co-organized the **3rd International Workshop on Cyber-Security Threats, Trust and Privacy Management in Software-defined and Virtualized Infrastructures** (SecSoft 2021) co-hosted at the 7th IEEE International

Copyright © Rainbow Consortium Partners 2020 -2022

Conference on Network Softwarization (NetSoft 2021) that was held from June 28 to July 2, 2021, in Tokyo, Japan.

- RAINBOW together with the EU-funded projects GUARD, SIMARGL, PALANTIR, INSPIRE-5Gplus, SIFIS-Home, ELECTRON and SDN-microSENSE, was co-organizer of the **4th International Workshop on Cyber-Security in Software-defined and Virtualized Infrastructures** (SecSoft 2022), hosted under the 8th IEEE International Conference on Network Softwarization (NetSoft 2022) in Milan, Italy from June 27 to July 1, 2022.
- On 4th of April, 2022 the Data & Web Science Laboratory of Aristotle University of Thessaloniki (AUTH) within the framework of its post-graduate studies program "Data & Web Science" and with the support of RAINBOW organized an online workshop on the topic of "**Processing Data in the Fog. The example of the RAINBOW Fog Computing platform**".
- On June 21, 2022 the RAINBOW consortium hosted an online workshop on **Edge Orchestration** with the participation of the Futurewei Technologies, member of the Linux Foundation Project Centaurus.

RAINBOW successfully organized a total of 6 webinars covering various aspects of its technological innovations as well as the industrial use cases that validated its solutions.

- Technical webinar "**Security and Trust in Fog Computing**"
  On February 22, 2022, IFAT with the support of DTU and POLITO, also involved in the security and trust for fog and cross-cloud services, organized a webinar to present the outcomes in the context of security and trust within the RAINBOW platform.
- Technical webinar "**The RAINBOW Data Management and Analytics Stack in Action!**"
  On May 27, 2022, UCY and AUTH partners organized a webinar showcasing a "hands-on" introduction to in-place data management and distributed data processing for fog computing with the use of the intuitive services comprising the RAINBOW Analytics Stack.
- Technical webinar "**Managing Service Level Objectives**"
  On June 15, 2022 the RAINBOW consortium hosted a technical webinar on managing Service Level Objectives (SLOs) and demonstrated the capabilities of the Polaris Framework, which is used to implement the SLO controllers and elasticity strategies in RAINBOW.
- Use case webinar "**Human-Robot Collaboration in Industrial Ecosystems**"
  On December 9, 2022, BIBA, the partner leading the use case on human-robot collaboration system in industrial ecosystems, hosted a live webinar on the challenges of this application and the use of RAINBOW's platform to ensure safe interaction in a shared workspace.
- Use case webinar "**Digital Transformation of Urban Mobility**"
  On December 6, 2022, CRF, the partner leading the use case on the digital transformation of urban mobility, together with partners LINKS and POLITO hosted a live webinar on the challenges and use of RAINBOW's platform in urban mobility scenarios.

Copyright © Rainbow Consortium Partners 2020 -2022

- Use case webinar "**Power Line Surveillance via Swarm of Drones**"
  On December 2, 2022, MSP, the partner leading the use case on power line surveillance via swarm of drones conducted a webinar summarizing the limitations and the challenges associated with drone missions and the solutions proposed by integrating drone systems with RAINBOW services.

RAINBOW has also organized a session hosted under the European Big Data Value Forum 2022. The session, which was co-organized together with the EU-funded cloud computing projects PLEDGER and MORPHEMIC, was titled "**Next-generation platforms for Europe's Cloud continuum**" and included a live discussion panel by prominent experts from academia and industry. The session highlighted the opportunities for the EU innovation landscape offered by the next-generation solutions developed by the 3 projects and illustrated how beyond the-state-of-the-art developments in the Edge, Fog and Cloud Continuum can strengthen EU competitiveness.

Other major dissemination actions that RAINBOW has undertaken during the project lifetime include:

Dissemination action	Description
Scientific and industrial events	In total RAINBOW has participated in **35** events with **18** of them including a project presentation. These attracted national, European and international audiences and included scientific conferences, workshops, webinars, community events, technology forums, industrial and innovation expos.
Clustering and collaboration with other initiatives	RAINBOW has established links and synergies with the **11** other projects and initiatives. This increased RAINBOW's visibility and enabled a series of cross-dissemination activities such as sharing of news, blog posts and project progress in social media; distribution of newsletters, knowledge exchange; participation and support at each other events; contributions to surveys and position papers, etc. In addition, RAINBOW has been an active member of the task force community established by the H-CLOUD and HUB4CLOUD CSA projects that aimed at supporting the European cloud computing ecosystem. Recently RAINBOW also joined the UNLOCK CEI community.
Joint activities	RAINBOW and its members have participated in **14** joint activities, such as scientific workshops, webinars, expert panels, surveys, position papers, undertaken together with other projects and initiatives.
Industry communities	By exploiting the collaborations and internal network of communications of the consortium members, news of

Copyright © Rainbow Consortium Partners 2020 -2022

Dissemination action	Description
	RAINBOW's progress and outcomes were disseminated to **11** national / regional communities. These include ICT developers and system integrators, technology SMEs, Cybersecurity community, start-ups, entrepreneurs and innovation communities, automotive industry, industrial production and logistics communities, aerial vehicles and flight systems community. RAINBOW also published about its platform in 3 corporate magazines aiming at raising awareness among the automotive, manufacturing and industrial production communities.

On the communication level RAINBOW has actively pursued the animation of several online channels to widen the reach and impact of its solutions.

Communication action	Description
Online presence	RAINBOW maintains a website containing lot of resources for the visitor. This includes: • All public deliverables of the project, publications, newsletters and training material. • A blog with over **100 posts** with several of them covering research challenges and technical applications on the domain of edge, fog and cloud computing. • A guide for developers on the installation and deployment for each released RAINBOW component. Up to the last month of the project the website had a total of **29,641 page views**, 6,821 unique visitors with the **average duration of visits** being **2min 17se**c. RAINBOW has been very active in social media and has established pages and accounts in LinkedIn, Twitter, Facebook, Research Gate, YouTube and Instagram in order to share all project news and activities of the project. The **total cumulative followers** are **1,193** while the **total number of posts** exceed **1,215**.
Communication material	RAINBOW published a total of **12 newsletters** covering the news, activities and progress of the project from the first year up to the completion of the project. RAINBOW produced **3 short videos** that provide an overview of the project, the use cases and its platform. In addition, has made available online all the **material and the**

Copyright © Rainbow Consortium Partners 2020 -2022

Communication action	Description
	recordings of the 6 webinars available at the project's YouTube channel https://www.youtube.com/@user-kw4op2xb3k/videos. RAINBOW has issued **13 press releases** to notify media on the achievement of project milestones or several important activities. RAINBW has created **6 promotional designs** for the project including leaflets, brochures and banners.

The reported accumulated reach of all digital and physical dissemination and communication activities is estimated at over **67,000**.

6.2.1 Scientific Publications

The strong research character of RAINBOW is made evident by the impressive number of 28 scientific publications that were achieved during the project and the mobilization even of non-academic partners in this direction. In overall, RAINBOW achieved the publication of 21 conference papers (with 3 more submitted during the last months of the project duration and awaiting acceptance) and 7 journal papers. Furthermore, three RAINBOW publications have received the best paper award in their category in the conferences they were submitted to.
The publications, which were in high-quality journals and international conferences organized by recognized scientific societies and associations, confirm that the various technologies and methods that were developed as part of the RAINBOW solutions advance beyond the-state-of-the-art. Scientific research results, such as those published by the consortium, have a positive impact in the diffusion and visibility of the project outcomes. Moreover, they offer opportunities for generating new knowledge, testing and obtaining feedback from the research community and even attracting further funding.

Table 6-1 RAINBOW Scientific Publications

Publications	Partner
6 conference papers	UCY
2 conference papers, 6 journal papers	POLITO
4 conference papers, 1 journal paper	AUTH
1 conference paper	UCY & AUTH
2 conference papers	DTU
2 conference papers	TUW
2 conference papers	IFAT

Copyright © Rainbow Consortium Partners 2020 -2022

Publications	Partner
1 conference paper	BIBA
1 conference paper	K3Y

6.2.2 Fogify

An indicative example of knowledge transfer has been the use of Fogify[7] from partner UCY; Fogify is an awarded open-source Fog Computing emulation framework which was employed in RAINBOW project and facilitated the evaluation of the use-cases and RAINBOW components. In a nutshell, Fogify is an award winning[8] framework that provides a toolset to model complex fog topologies which comprise of heterogeneous resources, network capabilities, and QoS criteria; deploy the modelled configuration and services using popular containerized infrastructure-as-code descriptions to a cloud or local environment; experiment, measure and evaluate the deployment by injecting faults and adapting the configuration at runtime to test different "what-if" scenarios that reveal the limitations of service before introduced to the public.

Design large-scale IoT testbeds
Fogify embeds a rich set of modelling abstractions where during the modelling phase users can describe the capabilities of their IoT services and the underlying infrastructure. Most importantly, no changes are required to the business logic of an application with the only assumptions that they adopt a micro-services architecture and that services are containerized using Docker. Hence, a typical workflow starts by annotating application services with infrastructure requirements and configuring network connectivity between infrastructure offerings.

Deploy IoT applications as a virtual lab
With an IoT testbed description in hand, users can take advantage of the Fogify SDK, to deploy the description in a virtual environment where the IoT application actually runs but with the infrastructure emulated on either the user's laptop or a computing cluster. This is possible thanks to the research backing Fogify where novel techniques have been designed for enforcing compute limits on host environments and extending overlay mesh with network rule chaining. The FogifySDK, as a python library, can be imported in notebook environments to interact at runtime with the Fogify Controller that is in charge of translating the Fogify model specification to underlying orchestration primitives. Fogify supports both Docker Swarm and Kubernetes in Docker (KIND).

Execute repeatable and reproducible "what-if" scenarios
Fogify users can create experiment scenarios and apply ad-hoc actions to deployed testbeds that will enable users to assess application performance and QoS. Actions supported by Fogify include a wide range of chaos testing tools, including changes to the

[7] https://ucy-linc-lab.github.io/fogify/
[8] Best Paper at IEEE/ACM IoTDi 2022 and Best demo at IEEE/ACM SEC 2020

Copyright © Rainbow Consortium Partners 2020 -2022

infrastructure, network QoS, injecting faults, vertical and horizontal scaling. The combination of monitored metrics with the analytical capabilities of the FogifySDK permits users to produce hidden insights about fog-oriented concepts such as Quality of Service, deployment cost and system reliability.

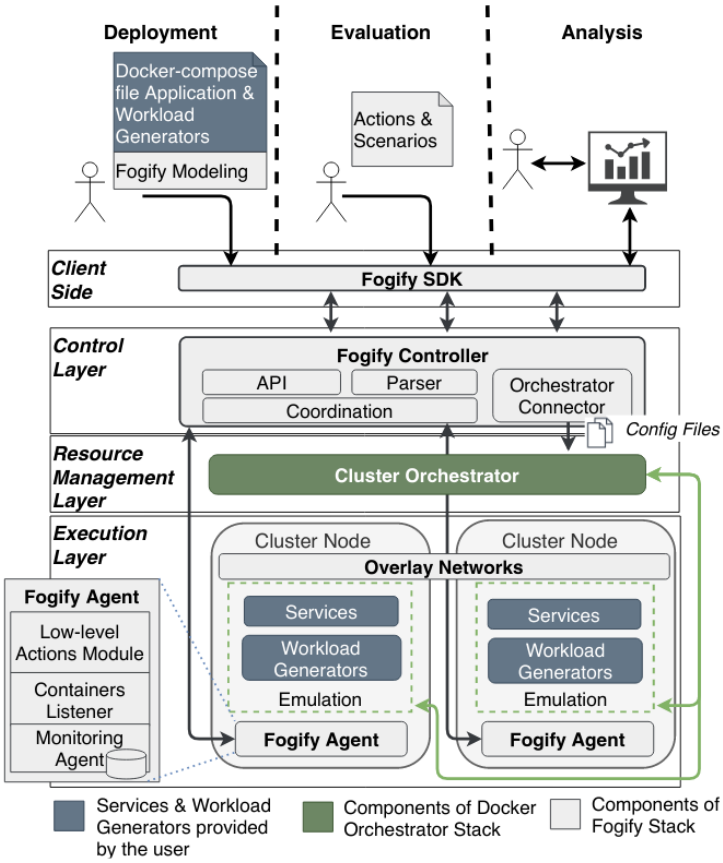

Figure 6-1 Fogify architecture

In summary, the Fogify emulation stack helps organizations in (i) designing large-scale IoT testbeds through modeling abstractions to configure cloud-edge nodes and networks; (ii) deploying the emulated infrastructure on a user's laptop or a cloud cluster for large-scale experiments; and (iii) rapidly defining reproducible experiments and "what-if" scenarios. The latter ensures that SMEs have accessibility to edge testbeds alleviating purchase and effort costs in infrastructure setup and rapid experiment design to test IoT applications' performance, energy footprint, and resilience in extreme conditions. Moreover, an organization can introduce Fogify in its DevOps pipelines of Fog applications enriching the testing phase with realistic Fog-related chaos testing and ensuring the performance and validity of an IoT application before the production deployment. Finally, Fogify could contribute to the research community since research organizations can utilize Fogify for large-scale experimentation or/and extend its ecosystem with new concepts, such as 5G networking (5G-slicer) or big data analytics (equality).

Copyright © Rainbow Consortium Partners 2020 -2022

Through the RAINBOW H2020 project, Fogify was used to test three very demanding pilot applications in the area of urban mobility, industrial manufacturing and autonomous drone swarms. Such pilots require wide-area testing over a plethora of devices while testbed design is a slow, labor-intensive and safety-critical process. All these hamper time-to-market and pose significant costs for SMEs. As an example, BIBA designs software and machinery for industrial IoT where human-robot collaboration are optimized while still obeying strict safety protocols. Testing optimization techniques during application prototyping is risky and can endanger personnel. However, through Fogify, developers design IoT testbed blueprints to configure their infrastructure needs and shared among colleagues and clients experiment scenarios to test various safety-critical functions and at the same time, assess real-time performance metrics. This is the innovation Fogify brings to the IoT landscape.

To increase the accessibility of the Fogify framework and, consequently, its success, we introduced systematic documentation management via a comprehensive documentation page and get-started guide for its use. Moreover, we published video presentations of research achievements of Fogify[9],[10] , conducted online webinars, and presented demos[11] [20] in multiple conferences [21] [22] providing the Fogify configurations and the deployed applications publicly[12]. Since we target to extend and capitalize on the achievements of the Fogify framework, we plan to implement a digital-twin extension on it by introducing the concept of cyber-physical emulation and providing the Fogify on demand (Emulation-as-a-service) by utilizing tools (such as terraform) or provider-specific functionalities (e.g., Amazon's AMI) as future exploitation and dissemination activities.

6.3 Hackathons

RAINBOW partners UCY and AUTH organized local hackathon activities with computer science students to allow the participants to experience first-hand the capabilities of the RAINBOW platform and utilize its tools to develop their unique solutions. The target audience was selected for their position as future ICT professionals and possible early adopters of next generation solutions such as RAINBOW's Fog Computing platform.

Hackathon on *"**Building optimized IoT applications for distributed environments**"*

On 3rd of December 2022 the Data & Web Science Laboratory of Aristotle University of Thessaloniki (AUTH), with the support of RAINBOW and within the framework of its Data & Web Science postgraduate studies program, organized a hackathon activity titled "Building optimized IoT applications for distributed environments". The activity was addressed to undergraduate and master students as well as PhD candidates of the

[9] Video of Fogify's presentation at SEC https://www.youtube.com/watch?v=d_6bJzJidas
[10] Video of Fogify's presentation at Dimokritos https://www.youtube.com/watch?v=z37I1QVFabg
[11] Fogify Demo Video https://www.youtube.com/watch?v=PthMM6rC89o
[12] Publicly available applications can be found in Fogify repo: https://github.com/UCY-LINC-LAB/fogify as well as in 5G-slicer repo: https://github.com/UCY-LINC-LAB/5G-Slicer

Copyright © Rainbow Consortium Partners 2020 -2022

Department of Informatics, Aristotle University of Thessaloniki. In total, 16 participants formed 5 groups of 2-4 members, which worked together as a team during the competition. The challenge called for the development of an IoT application based on a specified scenario and using the RAINBOW data management stack by exploiting the available fog resources in an optimal way so that the data exchange with the cloud services is kept to a minimum. The participants in an anonymous feedback session commented very positively on the user experience and usefulness of the RAINBOW platform for managing and deploying IoT applications based on this first-contact experience with it from a developer perspective.

Hackathon on ***"Building IoT applications with Docker containers"***

On the 3rd of December, 2022 the Laboratory for Internet Computing (LINC) of the University of Cyprus (UCY), under the context of RAINBOW, organized a Hackathon Tutorial event titled "Building Internet of Things Applications with Docker Containers". The Hackathon Tutorial was addressed to undergraduate and master students as well as PhD candidates, of the University of Cyprus, Department of Computer Science, with basic knowledge of coding (in Python) and Linux OS. The participants were asked to deploy a distributed Machine Learning inference application on a fog network of Raspberry PIs using Docker technologies. The main challenge was to deploy the latter application via the RAINBOW dashboard, extract user-defined metrics, and design a related analytics query. Participants benefitted from being introduced to the complete pipeline of the development and deployment of modern IoT applications in fog networks while getting themselves familiarized with the use of docker technology and with beyond the-state-of-the-art fog solutions such as the RAINBOW platform.

6.4 RAINBOW Adoption Guidelines

RAINBOW platform consists of various components which need to be installed and configured; in the following subsections a quick overview of the instructions for installing and using RAINBOW are described. More analytic presentation of those instructions may be found in deliverable D5.4 [4]. The following instructions are up to date at the time of writing this deliverable and are continuously updated based on our users' feedback at the ReadTheDocs page of RAINBOW: https://rainbow-h2020.readthedocs.io and at the RAINBOW GitLab public repository: https://gitlab.com/rainbow-project1/rainbow-installation.

6.4.1 RAINBOW Installation

RAINBOW installation is divided into three major stages. The first stage is the core platform setup that includes the docker engine, the Mesh Network, the Kubernetes cluster, and the RAINBOW components such as the Logically Centralized Orchestrator along with the Orchestration Lifecycle Manager and their subcomponents. The second stage is the Analytics Stack setup that includes the Monitoring, Data Storage and Analytic

Copyright © Rainbow Consortium Partners 2020 -2022

Services components. The third stage is the Dashboard setup that includes the Dashboard component.

RAINBOW supports the vast majority of Linux capable devices (i.e., VMs, wearables, drones, etc.) and the most widespread distributions, such as Ubuntu and Debian. The core platform along with the Analytics Stack are core elements and part of the master and worker nodes. The minimum execution requirements for the master node are at least 4 CPU cores, 8GB of RAM, 40GB of storage and x86 based CPU architecture and for the worker nodes are 2 CPU cores, 2GB of RAM, 20 GB of storage and either x86 or ARM based CPU architecture. For more advanced use cases, we propose a master node with at least 4 CPU cores, 16GB of RAM and worker nodes with 2 CPU cores and at least 4GB of RAM. Furthermore, we support GPU enabled devices, as also most of the devices which register under the /dev linux path. For the Dashboard component the minimum execution requirements are 2 CPU cores, 4GB of RAM, 20 GB of storage and x86 based CPU architecture.

The installation procedure is automated and provided in detail at the RAINBOW GitLab public repository: https://gitlab.com/rainbow-project1/rainbow-installation.

Core Platform setup

Name	Last commit	Last update
..		
calico.yaml	add installation script files	8 months ago
cjdns.txt	add installation script files	8 months ago
coreDnsDeployment.json	add installation script files	8 months ago
init-02-cjdns.sh	feat: add dashboard	5 months ago
init-03-cjdns-ipv6.py	add installation script files	8 months ago
init-03-cjdns-worker.py	add installation script files	8 months ago
init-04-configure-host.py	add installation script files	8 months ago
init-05-docker-debian.sh	add installation script files	8 months ago
init-05-docker-ubuntu.sh	add installation script files	8 months ago
init-06-docker-configure.sh	add installation script files	8 months ago
init-07-k8s-debian.sh	added missing update step after adding the kubernete...	4 months ago
init-07-k8s-ubuntu.sh	add installation script files	8 months ago
init-08-k8s-master.sh	add installation script files	8 months ago
init-09-k8s-master-configure.py	add installation script files	8 months ago
init-10-install-rainbow-orchestrator.sh	add installation script files	8 months ago
init-11-cjdns-master-credentials.py	add installation script files	8 months ago
init-12-k8s-join-master.py	add installation script files	8 months ago
k8s-init-config-ipv6.yaml	add installation script files	8 months ago
rainbow-v2-master.sh	update README and installation files	6 months ago
rainbow-v2-worker.sh	feat: worker script and README.md updates	4 months ago

Figure 6-2 Core platform installation scripts

Copyright © Rainbow Consortium Partners 2020 -2022

The first step of that stage is the installation and configuration of the master node that can be achieved by the execution of the *rainbow-v2-master.sh* of the provided installation scripts as depicted in Figure 6-2. This includes the installation of the docker engine, the prerequisites of the Mesh Network along with the Mesh Network itself, the Kubernetes Master and finally the RAINBOW components, the Logically Centralized Orchestrator along with the Orchestration Lifecycle Manager and their subcomponents. The only configuration that is required from the user is the necessary docker credentials for the involved containerized components.
After the successfully installation of the master component, the next step is the installation and configuration of the worker nodes that can be achieved by the execution of the *rainbow-v2-worker.sh* on each one of the worker nodes. This includes the installation and configuration of Mesh Network, the Node Level Agent (NLA) prerequisites as also the NLA itself. The script will configure the node and once it is ready, the NLA will continue by joining the node to the cluster automatically.

Analytics Stack setup

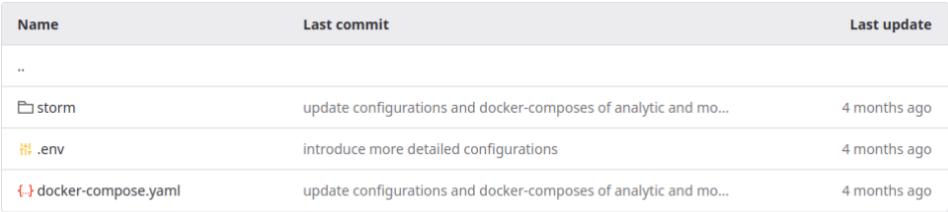

Name	Last commit	Last update
..		
storm	update configurations and docker-composes of analytic and mo...	4 months ago
.env	introduce more detailed configurations	4 months ago
docker-compose.yaml	update configurations and docker-composes of analytic and mo...	4 months ago

Figure 6-3 Analytics Stack setup – master node

The second stage is the installation of the Analytics Stack. The first step of that stage is the installation of the Analytics Stack on the master node. To complete this step, the user needs to configure the variables of the *.env* file that is depicted in Figure 6-3 with the variables that are shown in Table 6-2, and then execute the *docker-compose up* command.

Table 6-2 .env variables of Analytics Stack - master node

NODE_IPV6	Node's IPV6
NODE_IPV4	Node's IPV4
PROVIDER_HOSTS	The IPs of the nodes that the system will retrieve its data (all nodes' ips)
NODE_HOSTNAME	The hostname or ip of the node
STORM_NIMBUS_CONFIG_FILE	The path of Storm Nimbus configuration file
STORAGE_PLACEMENT	Enables and disables the placement algorithm of the storage. Default is False.
STORAGE_DATA_FOLDER	The folder that the data of storage will be stored

Copyright © Rainbow Consortium Partners 2020 -2022

Name	Last commit	Last update
..		
storm	feat: update README.md and add analytic-stack	5 months ago
.env	introduce more detailed configurations	4 months ago
docker-compose-arm32.yaml	fixes on data mngm node compose files	1 month ago
docker-compose-arm64.yaml	fixes on data mngm node compose files	1 month ago
docker-compose.yaml	fixes on data mngm node compose files	1 month ago

Figure 6-4 Analytics Stack setup – worker node

The second step is the installation of the Analytics Stack at the worker nodes. The procedure is similar with the previous step but as depicted in Figure 6-4, there are different docker-compose files in order to support the different architectures of the worker nodes.
Also, there are different environmental variables that need to be configured in the corresponding *.env* file that are presented and descripted in the following Table 6-3.

Table 6-3 .env variables of Analytics Stack – worker nodes

MONITORING_CONFIGURATION_FILE	The path of monitoring agent configuration file
STORAGE_RAINBOW_HEAD	Cluster head's IPV6
STORAGE_NODE_NAME	Node's hostname
STORAGE_PLACEMENT	Enables and disables the placement algorithm of the storage. Default is False.
STORAGE_DATA_FOLDER	The folder that the data of storage will be stored

Copyright © Rainbow Consortium Partners 2020 -2022

Dashboard setup

Name	Last commit	Last update
..		
.env	feat: rainbow metric exporter	2 months ago
README.md	feat: dashboard script fixes	4 months ago
cjdns.txt	feat: add dashboard	5 months ago
docker-compose.yaml	feat: rainbow metric exporter	2 months ago
init-02-cjdns.sh	feat: add dashboard	5 months ago
init-03-cjdns-ipv6.py	feat: dashboard script fixes	4 months ago
init-03-cjdns-worker.py	feat: add dashboard	5 months ago
init-04-configure-host.py	feat: add dashboard	5 months ago
init-05-docker-debian.sh	feat: add dashboard	5 months ago
init-05-docker-ubuntu.sh	feat: add dashboard	5 months ago
init-06-docker-configure.sh	feat: add dashboard	5 months ago
init-13-docker-compose.sh	feat: add dashboard	5 months ago
rainbow-dashboard.sh	feat: dashboard script fixes	4 months ago

Figure 6-5 Dashboard component scripts

The third and final stage of the RAINBOW installation is the Dashboard setup. The first step for that stage is the configuration and execution of *rainbow-dashboard.sh* scripts that includes the installation of the docker engine and the prerequisites of the Mesh Network along with the Mesh Network itself. The final step required the update of the *.env* file with the local IP address and the execution of the *docker-compose up* command.

6.4.2 RAINBOW Usage Guide

The on boarding to the RAINBOW platform starts from the Dashboard, where users can deploy and manage the applications through an easy-to-use user interface. A fully-fledged and frequently updated RAINBOW usage guide is available online as a ReadTheDocs page, at https://rainbow-h2020.readthedocs.io.

In this section, for the purpose of document completeness, we will present some of the basic parts of the Dashboard usage, by showing how a user can navigate, deploy and configure an application through it.
After successfully logging in, the landing page of RAINBOW is the main Dashboard view where an overview of the available resources along with monitoring information about the deployed applications are presented (Figure 6-6).

Copyright © Rainbow Consortium Partners 2020 -2022

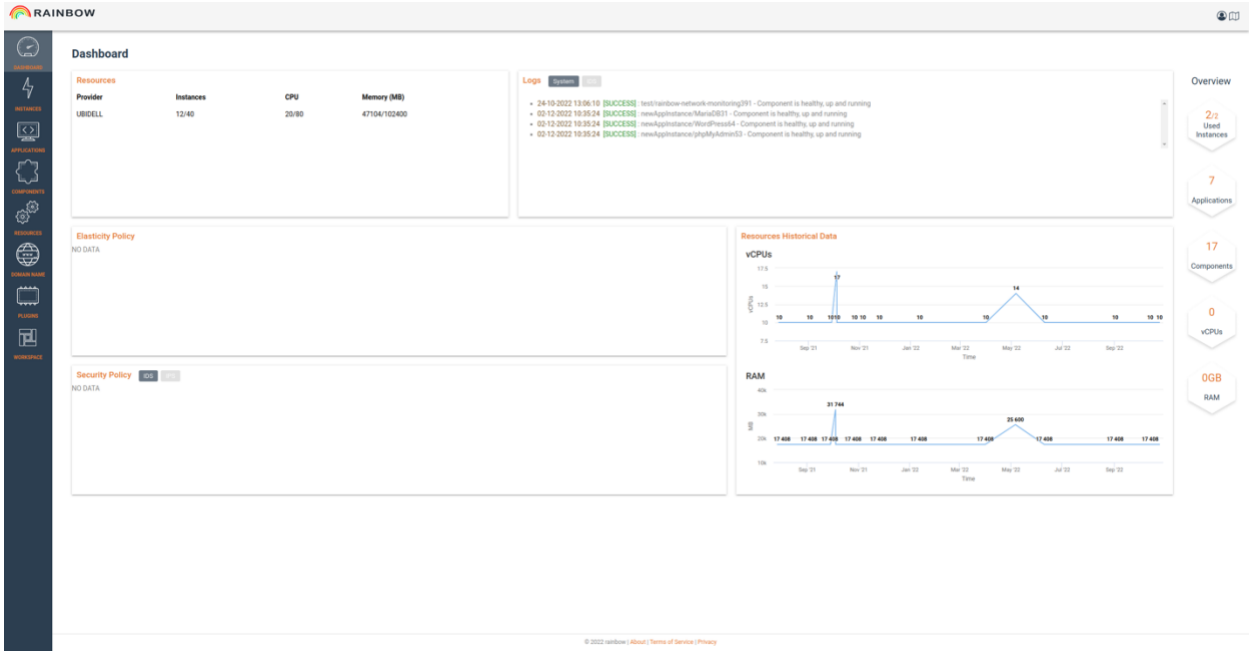

Figure 6-6 Main page of RAINBOW Dashboard

The user can view all the available components and create a new one or edit an existing one by configuring the necessary field with the desired values, as depicted in Figure 6-7 and Figure 6-8.

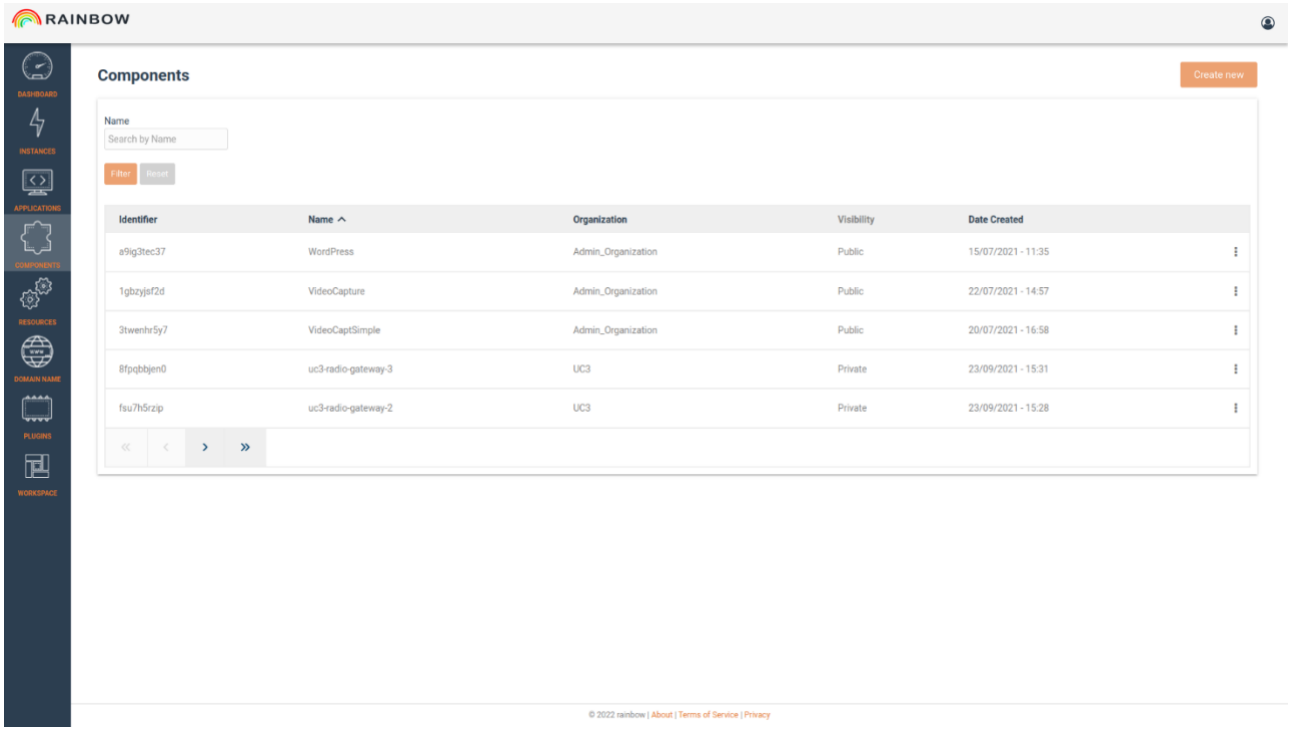

Figure 6-7 Components List in the RAINBOW Dashboard

Copyright © Rainbow Consortium Partners 2020 -2022

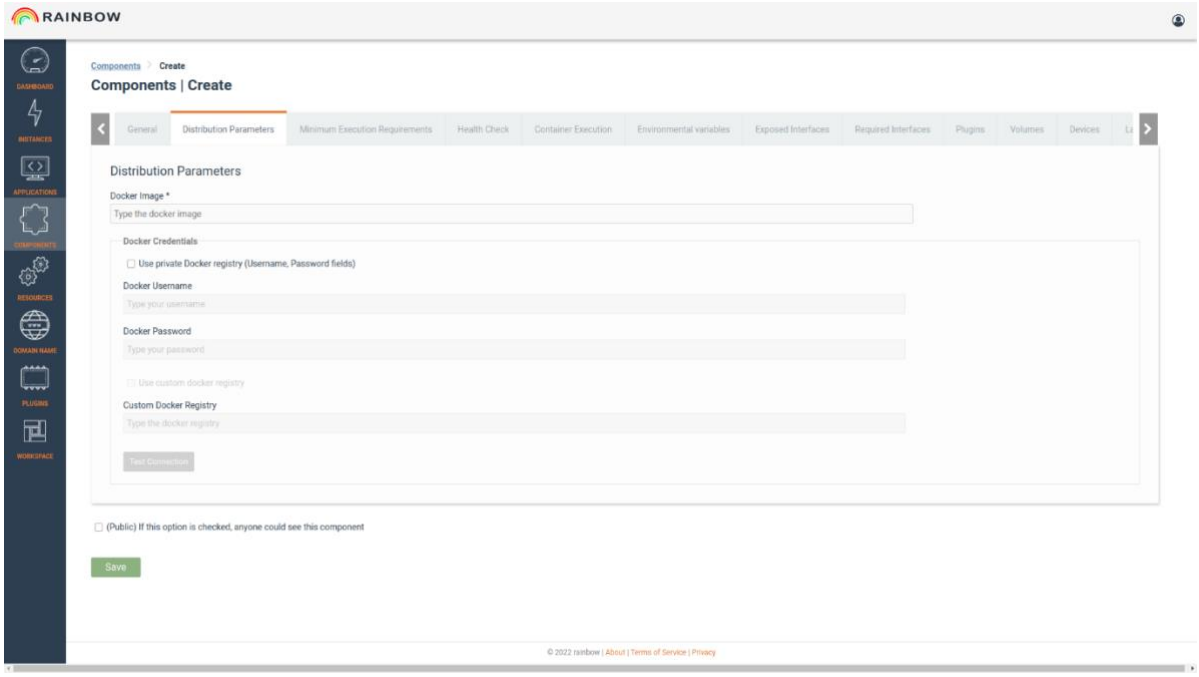

Figure 6-8 Component Configuration in the RAINBOW Dashboard

By using the defined components, the user can then graphically describe her/his application topology and store it as template through the Service Graph Editor (Figure 6-9).

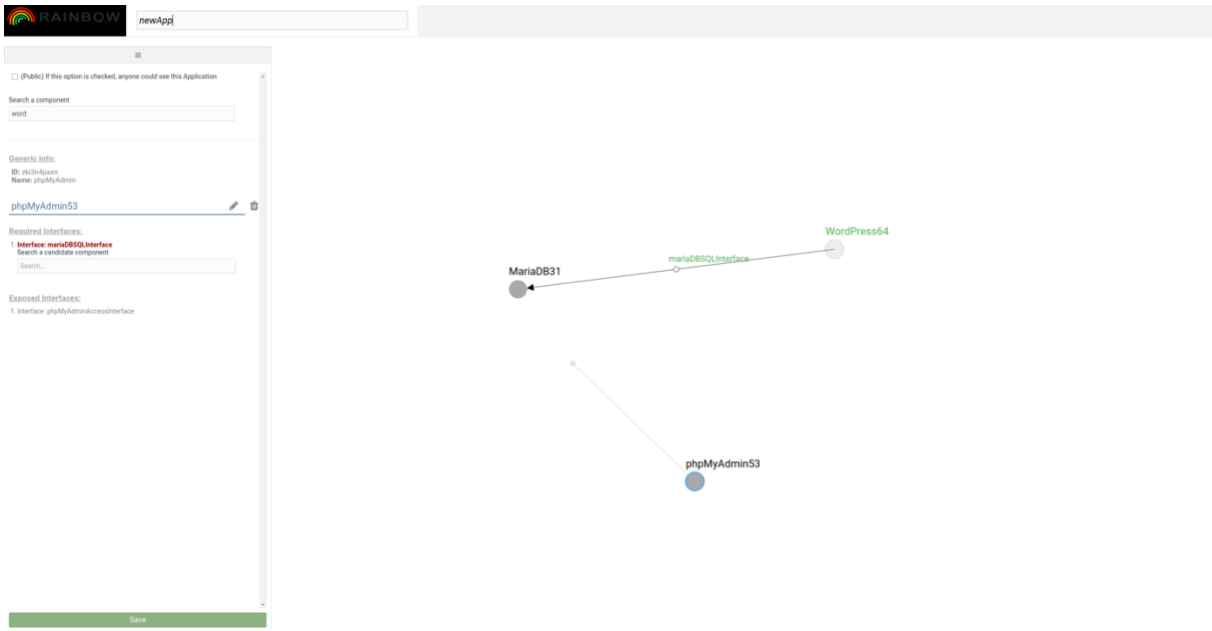

Figure 6-9 The Service Graph Editor of RAINBOW Dashboard

For an application to be deployed, the user must register an appropriate resource (Figure 6-10).

Copyright © Rainbow Consortium Partners 2020 -2022

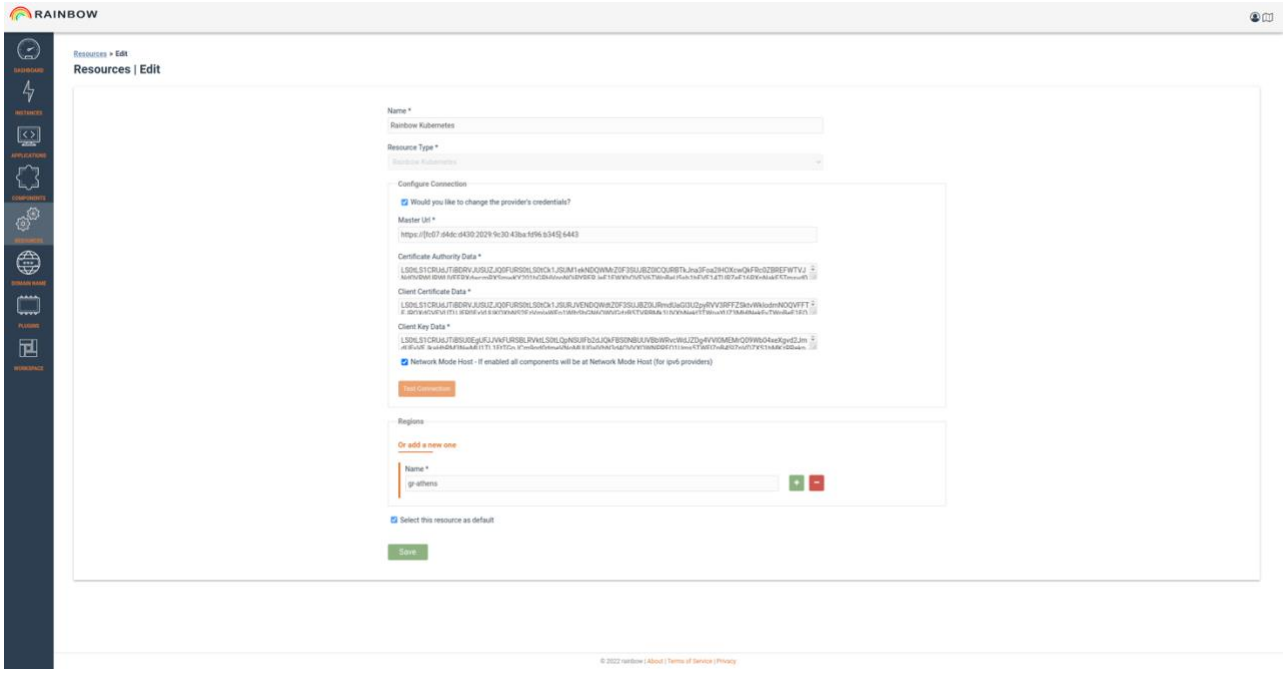

Figure 6-10 Resources management in the RAINBOW Dashboard

Following this, the application template can then be used to deploy it in the corresponding resource, as depicted in Figure 6-11.

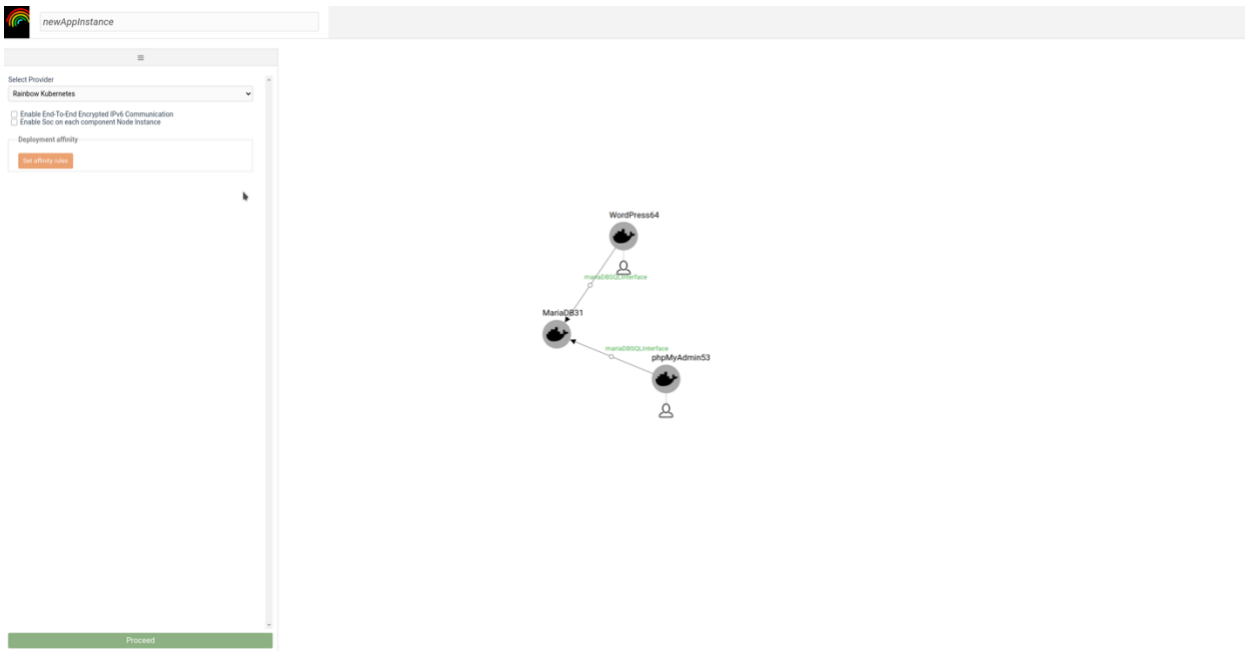

Figure 6-11 Application deployment procedure in the RAINBOW Dashboard

Once the application is deployed, it is listed in the application instances page (Figure 6-12).

Copyright © Rainbow Consortium Partners 2020 -2022

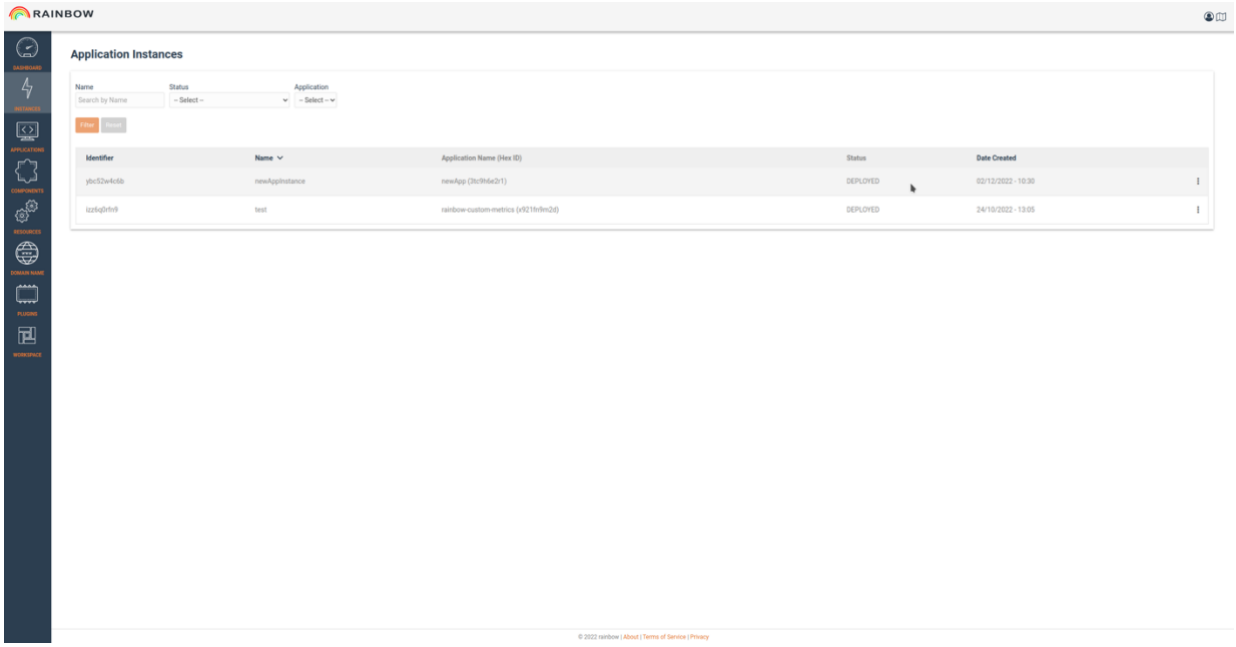

Figure 6-12 Application Instances list in the RAINBOW Dashboard

After the successful deployment of an application, the user can define analytics and elasticity policies by using the dedicated editors (Analytics Editor and SLO Editor respectively), as shown in Figure 6-13.

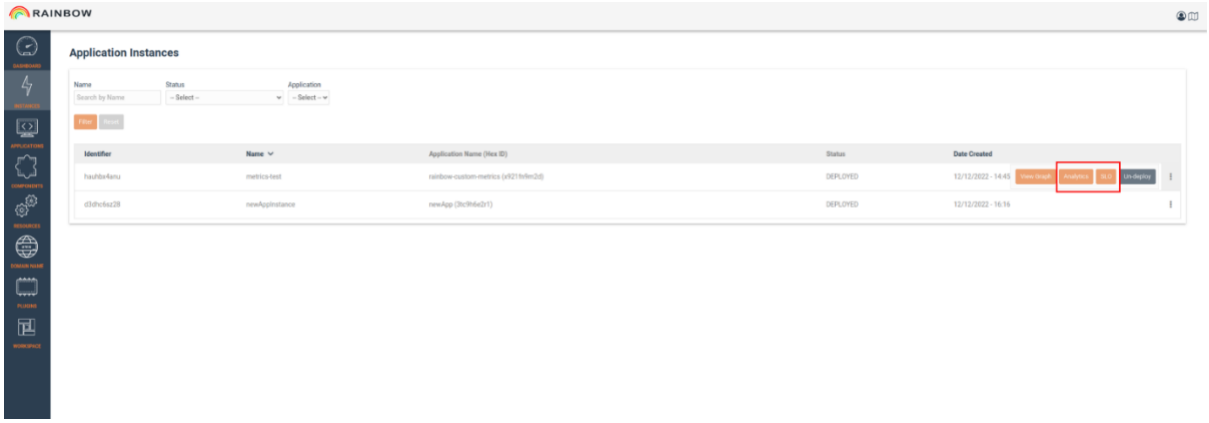

Figure 6-13 Analytics and SLO Editor in the RAINBOW Dashboard

In the Analytics Editor, the user can create the desired analytic metrics as depicted in Figure 6-14 and Figure 6-15.

Copyright © Rainbow Consortium Partners 2020 -2022

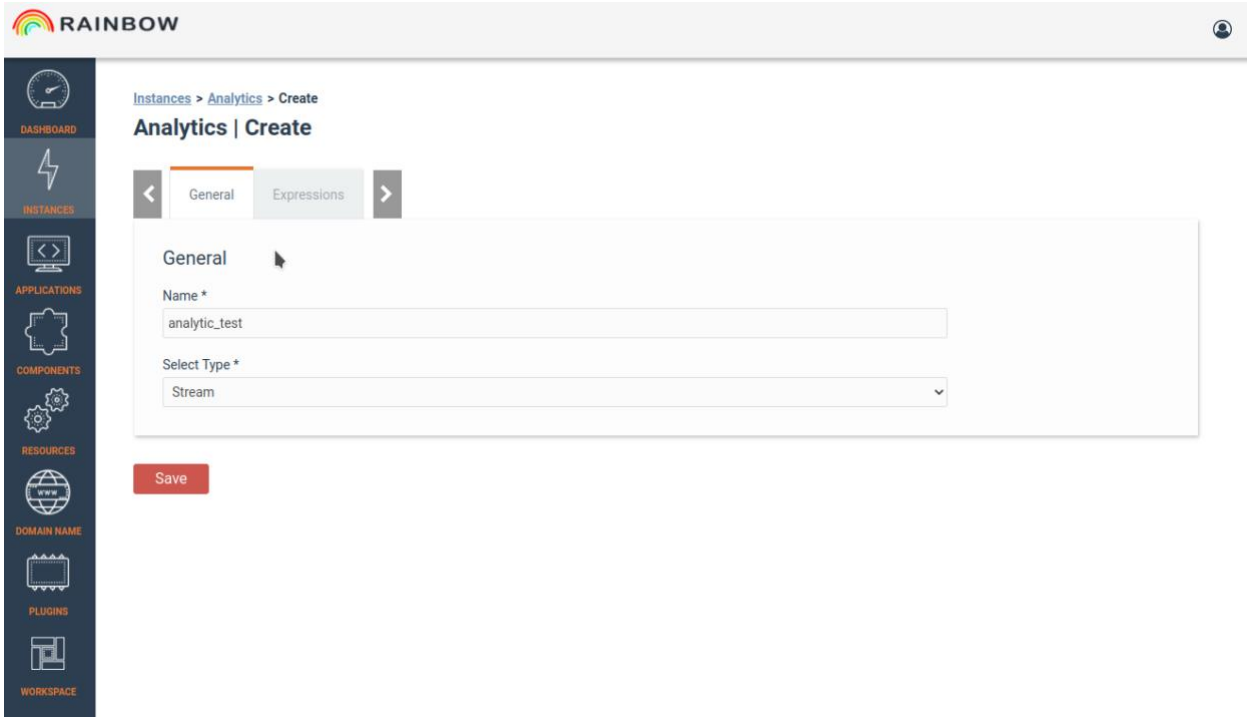

Figure 6-14 Creating analytics metric in the RAINBOW Dashboard

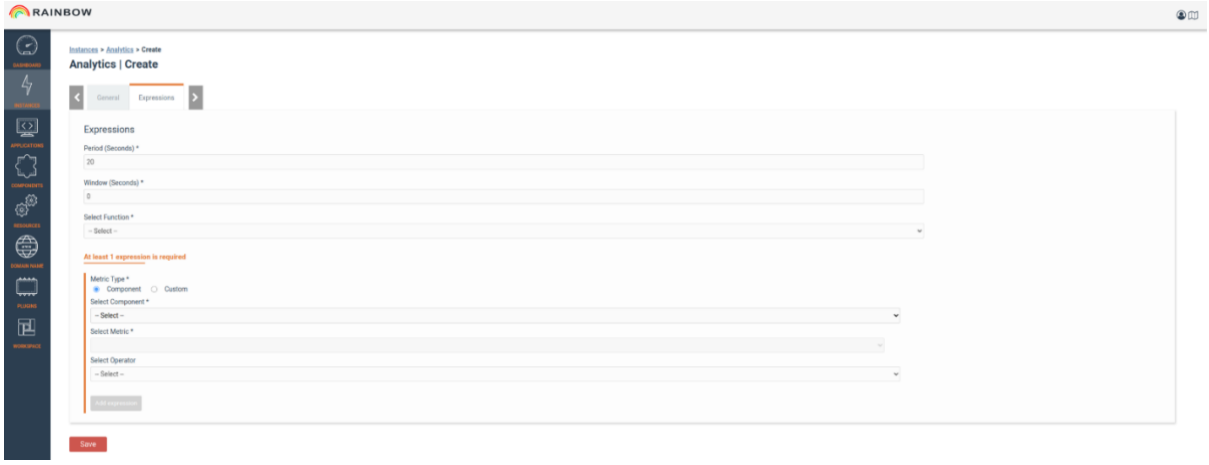

Figure 6-15 Adding expressions to the analytics metric in the RAINBOW Dashboard

SLOs can be added through the SLO Editor as depicted in the following Figure 6-16 up to Figure 6-19.

Copyright © Rainbow Consortium Partners 2020 -2022

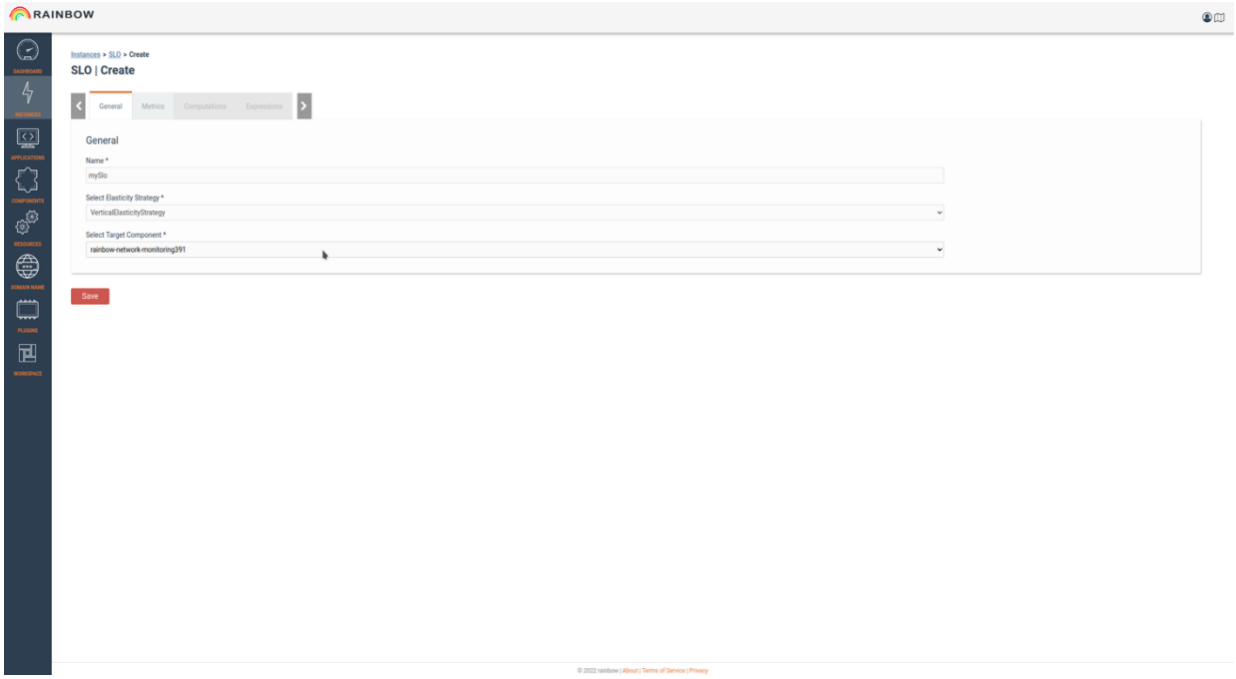

Figure 6-16 SLO creation in the RAINBOW Dashboard

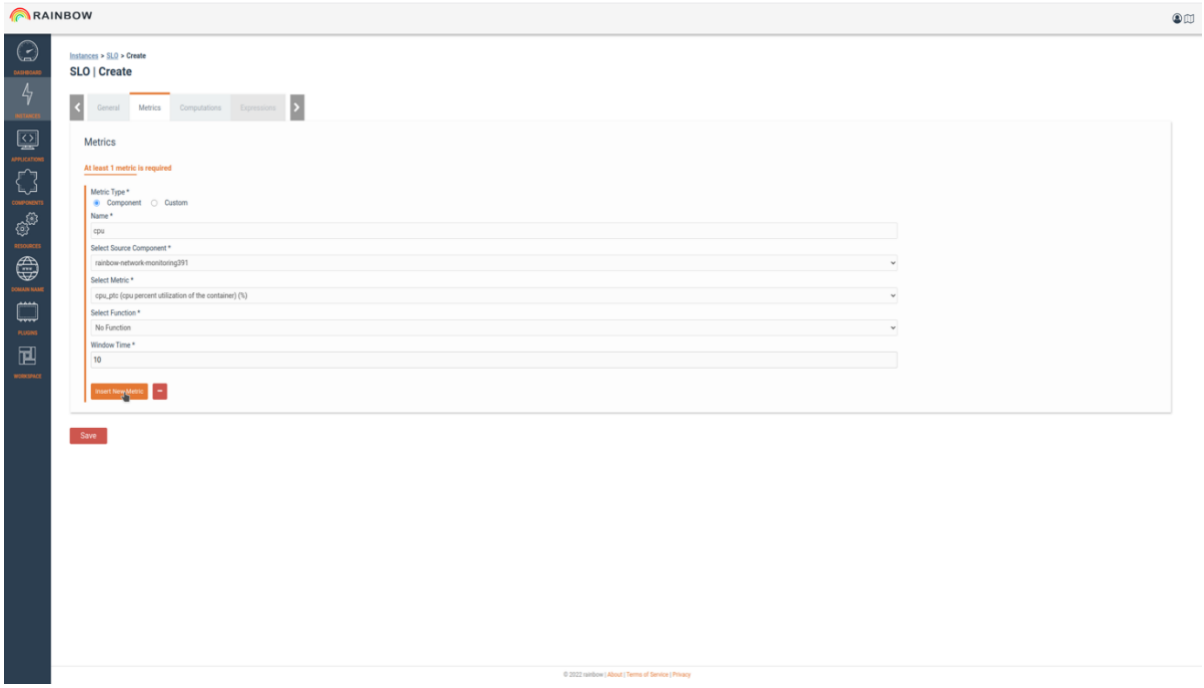

Figure 6-17 Adding metrics to the SLO in the RAINBOW Dashboard

Copyright © Rainbow Consortium Partners 2020 -2022

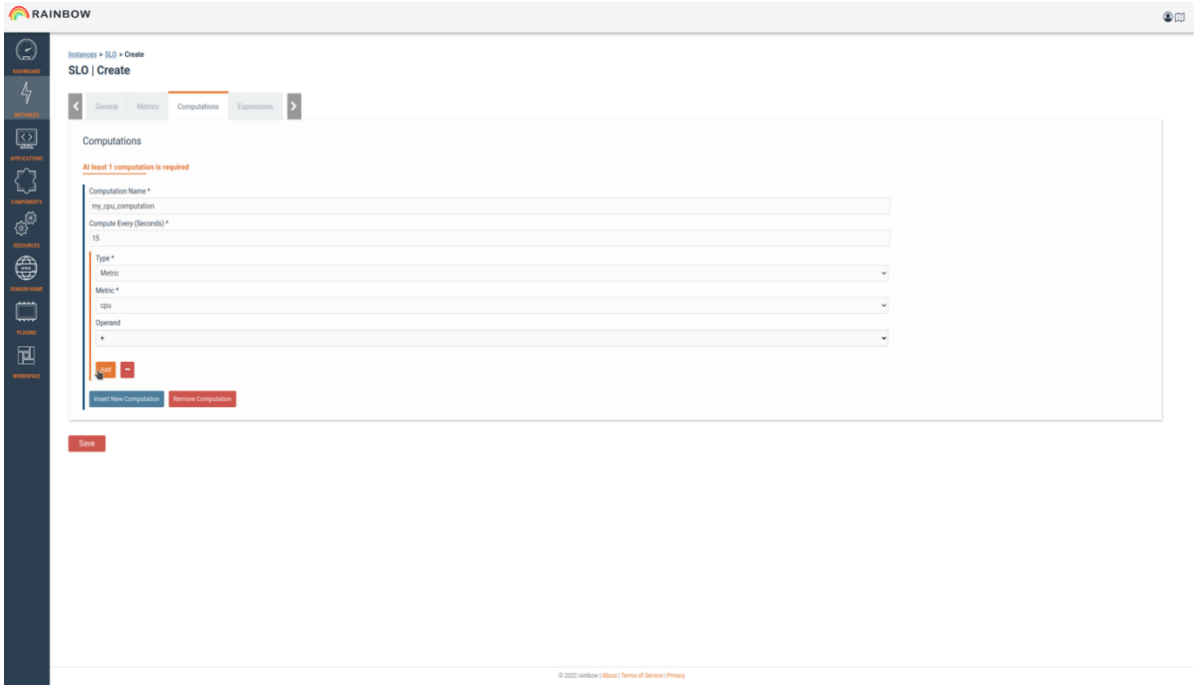

Figure 6-18 Adding computations to the SLO in the RAINBOW Dashboard

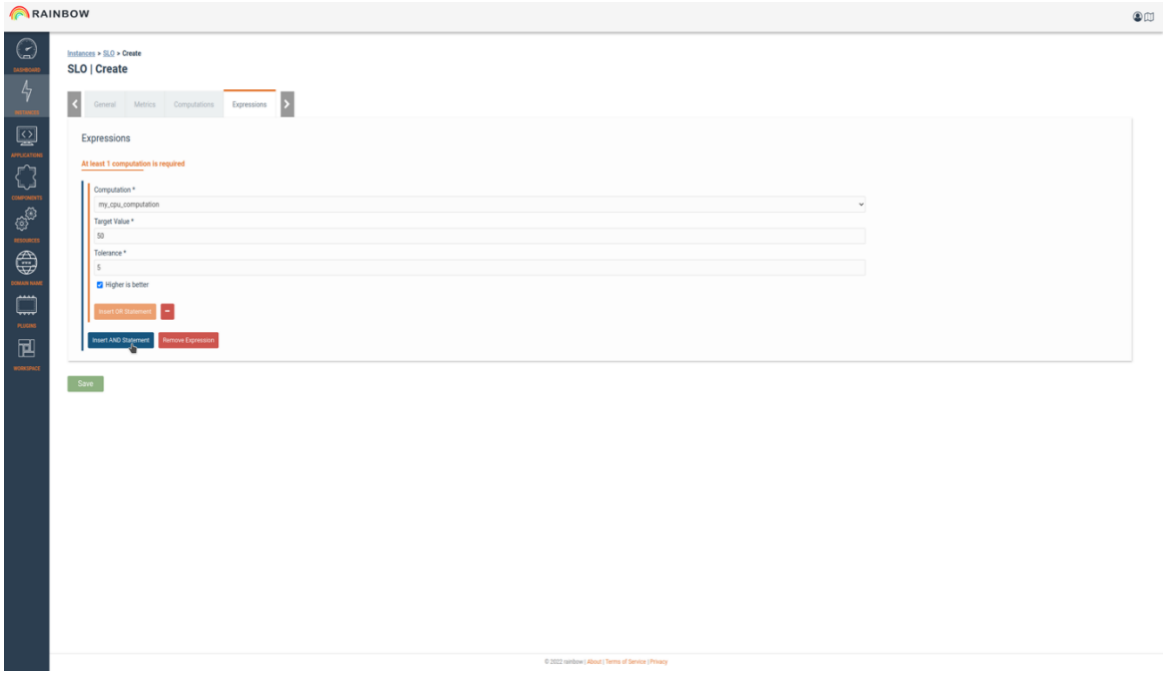

Figure 6-19 Adding Expressions to the SLO in the RAINBOW Dashboard

Copyright © Rainbow Consortium Partners 2020 -2022

7 Conclusions

The work under Work Package 7 "Dissemination, Exploitation and Communication" and, by extension, the RAINBOW project is completed with the current project impact assessment deliverable. It showcases the core impact caused by the RAINBOW project from all perspectives, from **technical output** and **scientific progress** to **contribution to standards** and policies, **strategic** and **economic initiatives**, to raising **awareness** and **public outreach**; all this throughout the course of the 36 months between 2020 and 2022, by gathering feedback from practically all the Work Packages of the project, from the point of view of both technical and demonstrator partners.

RAINBOW's goal was to let developers, service providers and infrastructure do what they do best without having to develop complex services for smart resource provisioning, monitoring and service orchestration. At the same time, developers, and IT solution providers (with emphasis on the SMEs that cannot invest on expensive proprietary modules) could get at their hands a complete solution which can guarantee device and mesh network security, data protection, identity management, anonymity, and resource integrity at all networking layers.

The above-mentioned goal has been achieved through the research and technical output of the project's work, in the form of RAINBOW platform components. As a whole, the RAINBOW platform contributes to the **strategic goals of Objective ICT-15-2019**: (i) development of an ecosystem that will respond to the future digitisation needs of industry and the public sector, (ii) assisting with the development of new cloud-based services and infrastructures in Europe and foster an industrial capability in the cloud computing sector, (iii) creating new opportunities to encourage European-based providers to develop and offer cloud-based services based on the most advanced technologies and (iv) leveraging research and innovation projects to support the development and deployment of innovative cloud-based services and next generation applications.

The first adopters of the RAINBOW platform have been the three advanced fog use cases, i.e., the three demonstrators described in the project. The provided RAINBOW functionalities were able to help add value to the operations, but at the same time feedback directed from the demonstrators back to the technical teams, sketched the advantages and the disadvantages the developed solution and led to its improvement. RAINBOW was initially built by the technical partners towards the direction set by the elicited stakeholder needs, namely: being **business-oriented** and **user-friendly**, offering **simplified deployment of complex applications** and easy-to-understand **visualisations of interdependencies** and **application architectures**. Working on from the early release and based on the received feedback, it has been successfully pivoted by the technical partners towards enabling **complete deployment automation** whilst ensuring **easy configurability**, supporting the creation of an **abstraction layer between user** and **infrastructure**, ensuring **elasticity** and **adaptability** in terms of hardware and node typologies.

Copyright © Rainbow Consortium Partners 2020 -2022

The advanced release of RAINBOW was set-up in three distinct operational environments, with deployments being performed on physical edges and testing and collection of metrics executed while the physical demo was operational. This has confirmed that **RAINBOW features a TRL of 7 "System prototype demonstration in operational environment"**.

It is of importance to also showcase the community benefits, from contributions to standardisation initiatives through the provision of open-source code (to the **Eclipse IoT community** and to **Linux Foundation**'s Centaurus project), to the open sourcing of RAINBOW services (the RAINBOW orchestrator, the monitoring agent, the analytic enabler, the storage agent, the overlay mesh networking stack, etc.), to the alignment with the architectural principles of OpenFog Consortium (rebranded as the **Industry IoT Consortium** in August 2021).

To conclude, transfer events held or endorsed by the RAINBOW consortium resulted in the sharing of know-how with interested stakeholders. The consortium made significant advancements in terms of the dissemination effort aiming at the wide dissemination of RAINBOW's results and focusing on the interaction with the target audience. In overall, RAINBOW organized **5 workshops**, attended or participated in **35 events**, performed **14 activities** jointly with other projects, and **delivered 28 scientific publications**, **6 webinars** and **14 open-source** beyond the state-of-the-art **components**. All the dissemination and communication activities whether these were performed digitally or in person have **an estimated reach of over 67,000**. Finally, through the organisation of Hackathon activities in partner universities, RAINBOW was able to receive further validation, as well as expand on the next generation of IT engineers.

Copyright © Rainbow Consortium Partners 2020 -2022

References

1 RAINBOW Consortium (2020) Description of Action
2 RAINBOW Consortium (2021) D5.2 RAINBOW Integrated Platform and Unified Dashboard – Early Release
3 RAINBOW Consortium (2022) D5.3 RAINBOW Integrated Platform and Unified Dashboard – Second Release
4 RAINBOW Consortium (2022) D5.4 RAINBOW Integrated Platform and Unified Dashboard – Final Release
5 RAINBOW Consortium (2022) D2.6 RAINBOW Secure Overlay Mesh Network - Final Version
6 RAINBOW Consortium (2022) D3.2 RAINBOW Orchestration Mechanisms
7 RAINBOW Consortium (2022) D4.2 Data Management Services
8 RAINBOW Consortium (2022) D7.5 Standardization Activities Report - Final
9 RAINBOW Consortium (2020) D1.1 RAINBOW Stakeholders Requirements Analysis
10 RAINBOW Consortium (2021) D1.3 RAINBOW Use-Cases Descriptions
11 RAINBOW Consortium (2022) D7.9 Market Analysis, Business Plan and Innovation Management - Final
12 RAINBOW Consortium (2022) D6.3 Human-Robot Collaboration Demonstrator – Final Release
13 RAINBOW Consortium (2022) D6.5 Digital Transformation of Urban Mobility Demonstrator – Final Release
14 RAINBOW Consortium (2022) D6.7 Power Line Surveillance Demonstrator – Final Release
15 RAINBOW Consortium (2021) D6.8 Validation Results and Performance Evaluation - Early Release
16 RAINBOW Consortium (2022) D6.9 Validation Results and Performance Evaluation - Final Release
17 RAINBOW Consortium (2022) D7.3 Open-Source Contributions, Dissemination, Clustering and Workshop Activities Report - Final
18 RAINBOW Consortium (2022) D2.4 RAINBOW Collective Attestation & Runtime Verification – Final Version
19 RAINBOW Consortium (2022) D7.7 Communication Activities Report - Final
20 M. Symeonides, Z. Georgiou, D. Trihinas, G. Pallis and M. D. Dikaiakos, "Demo: Emulating Geo-Distributed Fog Services," 2020 IEEE/ACM Symposium on Edge Computing (SEC), 2020, pp. 187-189, doi: 10.1109/SEC50012.2020.00031.
21 M. Symeonides, Z. Georgiou, D. Trihinas, G. Pallis and M. D. Dikaiakos, "Fogify: A Fog Computing Emulation Framework," 2020 IEEE/ACM Symposium on Edge Computing (SEC), 2020, pp. 42-54, doi: 10.1109/SEC50012.2020.00011.
22 [best paper award] M. Symeonides, D. Trihinas, G. Pallis, M. D. Dikaiakos, C. Psomas and I. Krikidis, "5G-Slicer: An emulator for mobile IoT applications deployed over 5G network slices," 2022 IEEE/ACM Seventh International Conference on Internet-of-Things Design and Implementation (IoTDI), 2022, pp. 115-127, doi: 10.1109/IoTDI54339.2022.00008.

Copyright © Rainbow Consortium Partners 2020 -2022