slices
DS

HORIZON 2020
H2020 - INFRADEV-2019-3

D2.5. Use cases validated

Acronym	SLICES-DS
Project Title	Scientific Large-scale Infrastructure for Computing/Communication Experimental Studies – Design Study
Grand Agreement	951850
Project Duration	24 Months (01/09/2020 – 31/08/2022)
Due Date	31 August 2022 (M24)
Submission Date	20 September 2022
Authors	Kostas Choumas (UTH), Nikos Makris (UTH), Christian Perez (Inria).
Reviewers	All partners

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 951850. The information, documentation and figures available in this deliverable, is written by the SLICES-DS project consortium and does not necessarily reflect the views of the European Commission. The European Commission is not responsible for any use that may be made of the information contained herein.

Scientific
Large-scale
Infrastructure
for Computing
Communication
Experimental
Studies
Design
Study

Executive Summary

This deliverable presents in detail the target use cases that the SLICES-RI architecture shall be able to support during its operation phase. The use cases have been selected based on the projected future demands and their purpose is to stress test the architecture and make extended use of the SLICES experimental components. This document presents the demanding KPIs per each use case, as well as the target KPIs for stressing the architecture and underlying testbed components.

In detail, the presented use cases are:

1) Rapid Resource Deployment for Physical Disaster Scenarios;
2) Smart-* applications: the smart cities' example;
3) Automated Construction and Demolition Waste Management using digital twin for buildings.

The use cases are described in detail regarding their scenarios of execution, as well as their deployment requirements for the SLICES-RI.

Table of contents

List of figures

Introduction

This document reports the reference use cases that will be used to validate the design space of SLICES. This work will evolve and evaluate in depth the reference architecture of SLICES-RI during the preparation phase in order to properly identify weather components of the architecture need to be re-engineered in order to meet the experimenters' demand. The use cases have been drafted by the consortium, based on the extended dissemination efforts of the project, the community and stakeholder engagement in the project, and try to sketch out applications that are currently in an incubation stage, but will be the norm beyond the next decade.

This deliverable provides the detailed description of three use cases that will be able to be executed, validated and reproduced over the SLICES infrastructure. All use cases deal with highly complex scenarios for Digital Science researchers, targeting three different scenarios with societal impact: physical disaster scenarios served by drones (Unmanned Aerial Vehicles or UAVs), a fully-automated smart city environment and a digital-twin assisted recycling model for buildings. All use cases will be made available to experimenters for reproducing and further extending them through SLICES.

1. Rapid Resource Deployment for Physical Disaster Scenarios

1.1. Use Case Description

In this use case, the fixed communication infrastructure could be destroyed or unavailable due to high workload demand. Furthermore, for rescue operations it could be critical to deploy additional on-demand computing and network resources at the proper place and time, in order to alleviate any remaining network infrastructure and collect data from remaining communicating devices such as mobile phones or sensors, towards helping to locate and rescue survivors. Mobile agents, such as robots or UAVs, are suitable for this kind of missions and can provide additional edge resources capable of processing the data at low latency and organizing the rescue operation.

To serve the survivor devices as much as possible, there is a need to predict the kind and amount of resources these devices will request and the location of these resources. Some mobile edge resources may need to be deployed sporadically and temporarily at different locations based on IoT devices needs and mobility. Thus, there is a need to anticipate the deployment of edge services and to estimate the time they will be required at a given place to decide whether it is worth deploying durable edge resources, or instead mobile temporary resources could suffice. In this latter case, the estimation of the location and quantity of required resources should be anticipated to allow their timely deployment. The deployment of edge resources will be such that a maximum of IoT devices can be served within the required latency, either directly or through multi-hop communications. Direct communications will be favoured for devices with very-low latency requirements, while multi-hop communications could be used for weaker latency requirements non-necessary communications.

The trajectory of distributed mobile edge devices should be consciously planned accordingly, taking in consideration the time restrictions (robots should be deployed at the proper place before we need them).

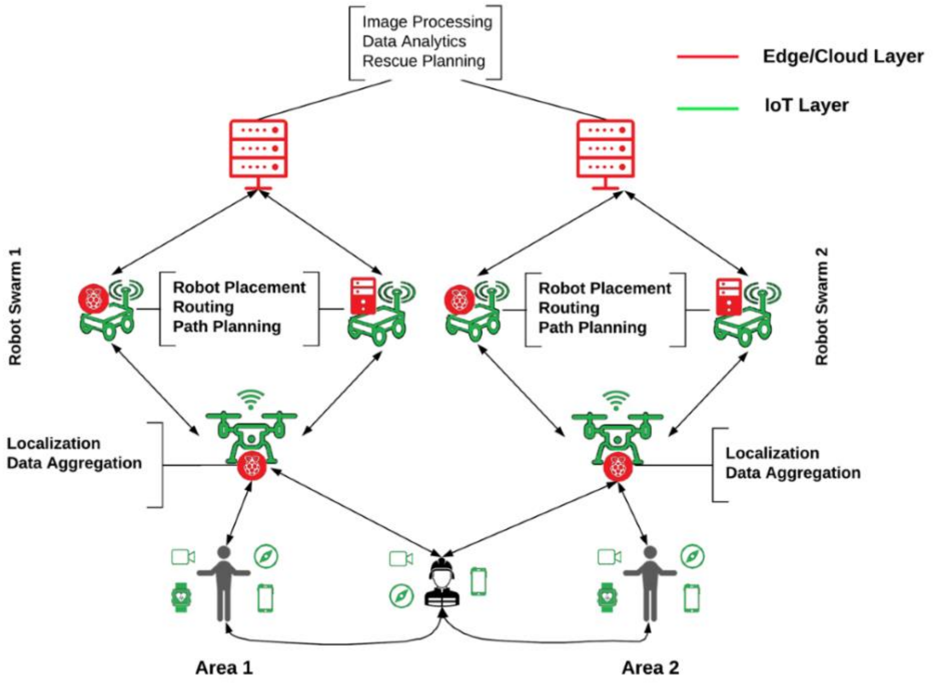

Figure 1: Network architecture for a physical disaster scenario that will be offered through SLICES

Using mobile devices will assist in eliminating the non-essential communications. Combining service differentiation and smart data offloading to UAVs or any other mobile device, there will be reduction of any unnecessary communication between users and overhead due to extensive signalling. Since the number of available robots may still be inadequate to serve all ground services, the prioritization of the applications, flows and devices is of paramount importance for the success of critical missions.

Under these circumstances, the priority will be given to the areas with many victims or of major importance for the completion of the mission, thus, the available swarms of mobile robots should be distributed accordingly. Even in the case of homogeneous mobile devices with identical computing and networking capabilities, the optimal allocation of them formulates a dynamic optimization problem, which depends on the size of the damaged area, the communication ranges, the propagation conditions, the data communication requirements (amount of data, frequency of collection, etc.) and the number and type of devices to serve. For example, if the victims are equally spread in different locations, the robot swarms would be equally scattered to deploy their resources in these areas so that the maximum number of devices would be served directly. On the other hand, the mobile robots will be driven to the most damaged area in order to serve the required network traffic.

This scenario illustrates the use and combination of the different control components of the framework, and in particular: 1) workload estimation in quantity, time and space, 2) resource allocation (tasks assignments to UAV and/or robots) and 3) path trajectory.

Existing platforms already allow the experimentation and testing of some components of such a use case, including robot self-deployments, data routing, distributed coordination, etc. but they lack the ability to provide edge and mobile edge services on one hand and the end-to-end experimentation of the complete use case in a single experimentation.

1.2. Use case operational KPIs

The KPIs used for the evaluation of the SLICES infrastructure on this use case are listed in the following table:

KPI	Description	Success
Accuracy on prediction of required resources	Machine Learning algorithms will be alleviated for predicting the workload in specific circumstances, based on the resources consumed until that moment. The percentage of the utilized resources, comparing to the estimated ones, will be the value of this KPI.	> 80%
Fair allocation of mobile devices	UAVs, robots and other mobile devices are fewer comparing to the edge resources. Their fair allocation is very critical and should be based on the necessities and the prioritization of the various rescues assisted by them. Given a specific set of communication necessities, without prioritizing some of them, the percentage of satisfying the requirements of each individual should be equal to the corresponding percentage of all other individuals. Hopefully, this percentage will be 100%, but in case that mobile devices are fewer, then all individuals should perceive the same percentage of utility. The difference between the minimum and maximum percentage will be the value of this KPI. Moreover, there will an opportunity to define individuals with higher priority, which would be able to perceive higher percentages of utility.	< 10%
Improvement on using mobile devices	Mobile devices should significantly improve the communication experience of the individuals, comparing to what the individuals would have experienced if the devices were not used. This KPI measures the improvement of using mobile devices as a percentage, comparing to the condition that they are not used. This KPI will be used for evaluating both the average improvement	> 20%

	over all individuals, as well as for evaluating the benefits for each individual.	

1.3. Use case KPIs for the SLICES architecture

This specific use case will stress test the architecture of SLICES by evaluating the following specific KPIs.

KPI	Description	Success
Network instantiation in an end-to-end manner	Time needed for instantiating the entire network using the SLICES reference architecture, enabling multiple geographical domains	Less than 1 minute
Time needed for processing data in real time	Data analytics and image processing needs to be executed in the cloud. Hence, proper resources (e.g., GPUs, HW accelerators and sufficient CPU cores and memory) should be available for supporting the processing requirements of the use case.	Processing lasts less than 10 seconds for even high frequencies of data collection
Edge-Core cloud communication latency	As decisions need to be taken in real-time, an efficient communication substrate needs to be established in an end-to-end manner. This KPI will evaluate the end-to-end time from all the collection locations (drones, remote cameras) to the core processing in the cloud.	< 5 msecs

2. Smart-* applications: the smart cities' example

2.1. Use Case Description

The number of sensors and actuators around us, projected up to 30 billion connected objects in the next few years[1], is a great opportunity to get a digital twin of the world and to design a set of helper software in our day-to-day life (or smart applications). However, if the IoT technology already starts to be commercialized, many research challenges have to be addressed to reach smart-* (or smart everything). Actually, this digital twin of the world is at the price of a huge amount of heterogeneous data to store, analyze and compute in various ways and by various software and services. This multi-dimensional heterogeneity (data type, computation type, software type etc.) is an important challenge, as well as the amount of data and services to manage, and as well as security and energy issues.

To illustrate these challenges, let us consider how a simple but usual event in a city such as a car accident can be managed tomorrow. First, it will be possible to detect this car accident by the sensors within cars, by users' smartphones, and video streaming information that could be can be analyzed (traffic light cameras, smartphones etc.). Such a car accident will initiate several notifications that should trigger new computations across different information systems (Police, Emergency services, Public transport, district/municipal services...). For example, police officers and vehicles should be re-scheduled and reorganized according to new incoming information. Video stream analyses should be started to automatically detect the number of vehicles or persons concerned by the accident, to deploy the adapted number of officers and the adapted types of vehicles. Also, a doctor could help predicting the type of injury caused by the accident, thus helping decisions to deploy emergency vehicles and persons. According to the type of injury, specificities of each hospital will be studied as well as the distance of hospitals to the accident. Moreover, decisions on doctors scheduling and GPS routing of emergency vehicles will be taken etc. Finally, the traffic should be adapted according to the accident and to the police management of the situation. GPS of users should be re-routed, and numerical simulations should be performed to predict the impact on the traffic jam, thus helping scheduling traffic lights etc. Of course, if a second accident happens while the first one is not ended, new decisions should be taken in order to handle both accidents according to police, hospital and traffic services.

Similarly, to the disaster scenario, there is a few testbeds that allow researchers to investigate some of the challenges of car-accident use-case. For example, distributed decision support system can be evaluated on top of testbed such as Grid'5000 in Europe. However, there is no testbed that enables the validation of such a use-case overall. That is, a testbed that allows researchers to execute representative scenarios in a controlled environment. Concretely, SLICES will help researchers observe and analyze all interactions between the IoT devices and the cloud resources where the applications are executed. It will help address the sizing challenge of edge resources in terms of computation, storage and network needs. Moreover, thanks to SLICES, researchers would be allowed to evaluate new distributed system/network building blocks such as publish/subscribe systems that can cope with the scalability and the velocity of such smart-cities environments (just imagine a large city such as Paris or Milan). Such building blocks are critical for the advent of smart-* applications.

---

[1] "State of the Market: Internet of Things 2016", https://www.verizon.com/about/sites/default/files/state-of-the-internet-of-things-market-report-2016.pdf, [Last accessed: 31 August 2022].

2.2. Use case operational KPIs

The KPIs used for the evaluation of the use case are listed in the following table:

KPI	Description	Success
Time needed for event handle	The time needed for the detection of an event in a smart-city environment, such as an accident, as well as the duration of the corresponding decision-making process, such as the trigger of multiple responses from various systems, including Police, Emergency services, Public transport, district/municipal services.	< 1 sec
Accuracy on successful event detection	The efficient utilization of the sensors and actuators in predicting the events of our interest, is measured using the recall score, which is calculated by dividing the true positives (TP) by the true positives (TP) plus the false negatives (FN): TP/(TP + FN).	> 95%

2.3. Use case KPIs for the SLICES architecture

This specific use case will stress test the architecture of SLICES by evaluating the following specific KPIs.

KPI	Description	Success
Network instantiation in an end-to-end manner	Time needed for instantiating the entire network using the SLICES reference architecture, enabling multiple geographical domains (e.g., different cities)	Less than 1 minute
Time needed for processing data in real time	Data analytics and deep data processing executed at the network edge for supporting the different smart city applications.	Processing lasts less than 1 second at the edge
Edge communication latency	As decisions need to be taken in real-time from the network edge, an efficient communication substrate needs to be established, relaying data to the cloud only when needed.	< 5 msecs

3. Automated Construction and Demolition Waste Management using digital twin for buildings

3.1. Use case Description

The construction industry is the main consumer of mineral and other non-renewable resources in the EU and generated about 35.7% of the total waste in the European Union in 2018, most of which was produced during the construction and demolition stages. Although there have been a number of initiatives to tackle issues of Construction and Demolition Waste (CDW), landfilling or downcycling of CDW is still prevalent. There are two broad reasons for the current unsatisfactory situation regarding CDW: low rates of recovery in a few EU countries and low value of reused or recycled materials/products. There is generally a lack of confidence of end-customers in the quality of recycled CDW due to improper sorting and contamination and low control over the material quality and homogeneity. First, the volume of materials entering CDW stream has to be minimized and, secondly, CDW has to be treated and better controlled to provide materials for reuse of high value products, with traceability and quality assurance.

Waste traceability is one of the biggest challenges in the construction industry. Digital approaches are key to facilitate managing information through the various stages of the built asset, while improving the circular economy in many situations with the use of technologies such as Blockchain or Artificial Intelligence. In this use case, a digital twin will be established, that will integrate the different stages of CDW management, by providing an integrated digital approach to enable waste traceability and management, where built asset project information is managed through the whole life cycle. Such an approach will establish a standardised information management workflow/process. This in turn, will serve to integrate the key stakeholders, ensuring that the right information is provided and accessible by the right ones, at the right time, providing a single view of the truth of the built asset information model at key stages in the life cycle (i.e., design, construction, operations, deconstruction, new products life), to effectively enable CDW management. The digital twin will be established through the adoption and customisation of a cloud-based collaboration solution, underpinned by existing standards, protocols, etc. for information management of digital twins, etc. This will facilitate a seamless flow of the required built asset information through the whole-life and CDW management process enabled through the interoperability of the different tools through open exchange standards. The required aspects of information security, privacy, and confidentiality will also be considered in alignment with relevant standards. Analytics and smart tools will be integrated to demonstrate the real-time dashboarding and reporting of the CDW management information to support more effective decision making. In addition, an automated integration of the tools for CDW management with the digital twin will be explored, in order to further demonstrate the enhancement of the seamless flow of information throughout the process. The operation of the digital twin will also be validated through the implementation of the use case.

Figure 2: Building digital twin

3.2. Use Case Operational KPIs

The KPIs used for the evaluation of the SLICES infrastructure on this use case are listed in the following table:

KPI	Description	Success
Reduction in time needed for estimation of produced waste	The digital twin developed for each building will enable the quick and accurate estimation of the building materials, as well as their quantities. The utilization of this digital twin, which will be updated automatically during the whole life of the building, will reduce significantly the time needed for estimating the waste produced by the demolition of this building.	> 70%
Waste reduction percentage	The digital twin will facilitate the reconstruction of building damages, precisely estimating the quantities and the materials needed for this reconstruction. It will reduce the ordered quantities of e.g. concrete, ordering not too far from what is needed and avoiding the waste production from materials that remained unused at the end of the reconstruction process.	> 30%

3.3. Use case KPIs for the SLICES architecture

This specific use case will stress test the architecture of SLICES by evaluating the following specific KPIs.

KPI	Description	Success
Network instantiation in an end-to-end manner	Time needed for instantiating the entire network using the SLICES reference architecture, enabling multiple geographical domains (e.g. different cities)	Less than 1 minute
Time needed for VNF deployment	The time required for the VNF deployment.	< 1sec

Digital twin communication latency	Latency for communicating real-time measurements with the network edge, feeding into the digital twin operation.	< 4 msecs
Time needed for processing data in real time	Data analytics and deep data processing executed at the network edge for supporting the digital twin operation (non-realtime).	Processing lasts less than 1 second at the edge

4. Conclusion

The three use cases will be exploited to thoroughly evaluate the SLICES infrastructure. The physical disaster scenarios served by the UAVs, the fully-automated smart city environment and the digital-twin assisted recycling model for the buildings are related to challenging problems with huge societal and environmental impact. The utilization of the SLICES infrastructure for building innovative and efficient solutions for these issues, provides sufficient guarantees for the great capabilities of the SLICES infrastructure and its potential to be exploited in various scenarios with great diversity. The defined KPIs will be used to precisely measure the improvements succeeded in each use case, the benefits of utilizing the SLICES infrastructure comparing to what happened before. Of course, during the long duration of SLICES-DS, the KPIs will be elaborated and enhanced even more, based on the experience gained from the project evolution. In any case, the given KPIs provide a minimum set of indicators for evaluating the success of the SLICES infrastructure.

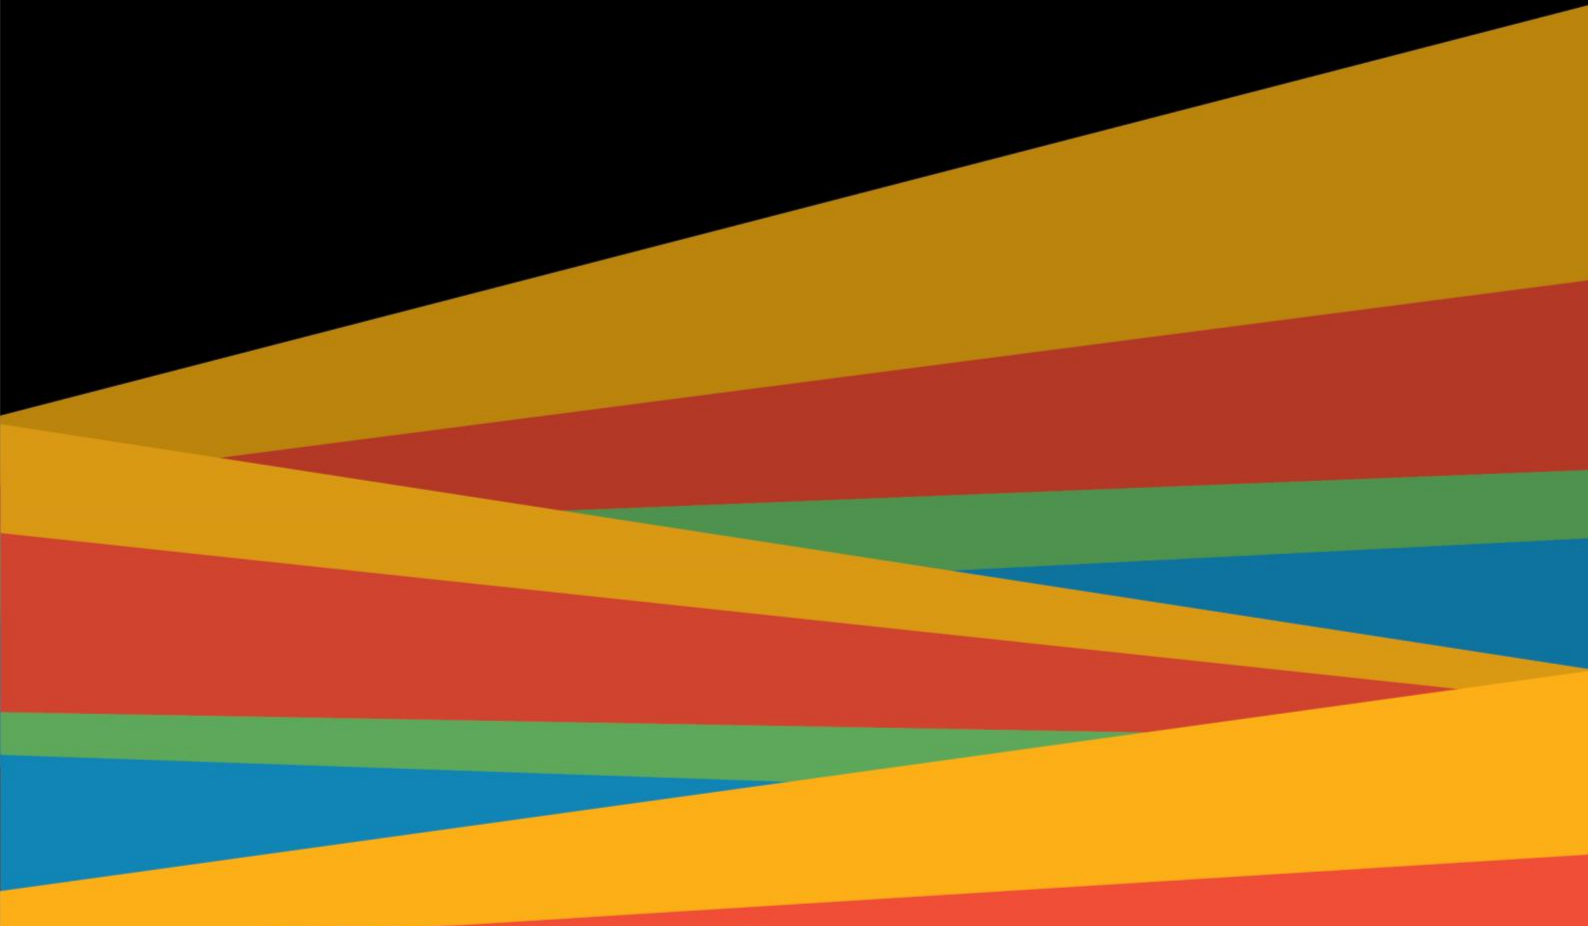

slicesDS