Grant Agreement No.: 731677
Call: H2020-ICT-2016-2017
Topic: ICT-13-2016
Type of action: RIA

D5.8: FLAME Replication Process v2

August Betzler, Carolina Fernandez (i2CAT)

Sebastian Robitzsch (IDE)

Navid Solhjoo, Rafael Guimaraes (University of Bristol)

Nishanth Sastry (King's College London)

Paolo di Francesco (Level7)

Stephen C. Phillips (IT Innovation)

30 April 2020

This document builds on top of D5.1 "FLAME Replication Process", being the second and final version of the FLAME replication process documentation. It provides a comprehensive guide on how to replicate FLAME in cities, focussing on the technical aspects, whereas the business aspects of the replication process described in D5.1 still stand. The installation and operation of FLAME in the replicator cities has led not only to the definition of very specific workflows and procedures for the different roles (infrastructure provider, platform provider, etc.) to validate the operability of FLAME, it has also resulted in the development in a series of tools designed for that purpose. Finally, we also describe the preparatory steps and dedicated workflows that have to be followed by experimenters who want to make use of FLAME.

Work package	WP 5
Task	Task 5.4
Due date	30/04/2020
Submission date	30/04/2020
Deliverable lead	I2CAT Foundation
Version	Draft for technical review: 10/11/2019; v1: 30/04/2020
Authors	August Betzler (i2CAT Foundation); Sebastian Robitzsch (IDE); Navid Solhjoo, Rafael Guimaraes (University of Bristol); Stephen C. Phillips (ITInnov); Paolo Di Francesco (Level7), Nishanth Sastry (KCL)
Reviewers	Steven Poulakos (DRZ)
Keywords	Replication, Platform, Guidelines, Infrastructure, SDNs, Requirements

DISCLAIMER

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 731677.

Project co-funded by the European Commission in the H2020 Programme		
Nature of the deliverable:	**R**	
Dissemination Level		
PU	Public, fully open, e.g. web	✓
CL	Classified, information as referred to in Commission Decision 2001/844/EC	
CO	Confidential to FLAME project and Commission Services	

Co-funded by the Horizon 2020 Framework Programme of the European Union

EXECUTIVE SUMMARY

This document serves in first instance as a manual/guideline that describes which steps have to be taken in order to replicate FLAME in an infrastructure and how a deployment can be validated. It is a formal write-up of the best current practice information made available to the FLAME open call 2 replicators. This includes specific guidelines for the infrastructure providers that have to make sure their infrastructure fulfils the requirements. Also, the use of tools developed for the purpose of an automated evaluation is detailed here. Further, it gathers valuable insights obtained from the cities in which FLAME has been deployed and also provides an overview of the replication status in the new replicator cities. Beyond that, a series of suggestions and recommendations are made for replicators that originate from the experience gained during the deployment and operation of the FLAME platform across the cities. Finally, this document also presents a workflow that experimenters have to follow in order to be ready to deploy their experiments in a FLAME-enabled testbed.

Please note that a draft version of this document was submitted in November 2019. To keep track of the changes between the draft and this final version, we include a changelog that lists all additions and modifications made to the document (Section 9.6).

Co-funded by the Horizon 2020
Framework Programme of the European Union

TABLE OF CONTENTS

Co-funded by the Horizon 2020 Framework Programme of the European Union

Co-funded by the Horizon 2020
Framework Programme of the European Union

LIST OF FIGURES

Co-funded by the Horizon 2020
Framework Programme of the European Union

LIST OF TABLES

Co-funded by the Horizon 2020
Framework Programme of the European Union

ABBREVIATIONS

AP	Access Point
API	Application Programming Interface
AZ	Availability Zone
CLMC	Cross Layer Management and Control
COTS	Commercial Off-the-Shelf
CPU	Central Processing Unit
DC	Data Centre
DHCP	Dynamic Host Configuration Protocol
FLIPS	Flexible IP Services
FQDN	Fully Qualified Domain Name
HOT	HEAT orchestration template
IEEE	Institute of Electrical and Electronics Engineering
IP	Internet Protocol
KPI	Key Performance Indicator
LAG	Link Aggregate Ports
LAN	Local Area Network
MAC	Media Access Control
MGMT	Management
MIMO	Multiple-Input Multiple-Output
MPLS	Multi-protocol Label Switching
NFV	Network Functions Virtualization
NIC	Network Interface Controller
QoE	Quality of Experience
QR	Quick Response
RAM	Random Access Memory
RAN	Radio Access Network

Co-funded by the Horizon 2020
Framework Programme of the European Union

REST	Representational State Transfer
SDN	Software Defined Networking
SFC	Service Function Chain
SR	Service Router
SSH	Secure Shell
SSID	Service Set Identifier
TCAM	Ternary Content-Addressable Memory
TCP	Transmission Control Protocol
TOSCA	Topology and Orchestration Specification for Cloud Applications
UE	User Equipment
UoB	University of Bristol
VLAN	Virtual Local Area Network
VM	Virtual Machine
VNF	Virtual Network Function
vSwitch	Virtual Switch
VXLAN	Virtual Extensible LAN
Wi-Fi	Wireless Fidelity

Co-funded by the Horizon 2020
Framework Programme of the European Union

1 INTRODUCTION

With FLAME now being deployed in 4 locations (Bristol, Barcelona, Kings College London (KCL), Sicily (Buseto Palizzolo)) it has been proven that FLAME can be replicated in a variety of different locations. Each of the current deployments can be considered unique in many different aspects, such as the infrastructure (compute, radio access, networking, etc.), type and extension of public spaces covered and the number and roles of the personnel participating in the project, to name a few. The four locations have one thing in common: they satisfy a set of requirements that need to be fulfilled wherever FLAME is to be operated. In Section 2 of this deliverable these requirements are clearly stated: the necessary software, hardware and infrastructure elements are determined. Further, based on the experience gained from the replicators, deployment guidelines are provided and the existing configurations are presented. Moreover, to cope with the heterogeneity of each replicator, the FLAME project has developed not only a series of guidelines and best practices for replicators that should be followed, but it also has produced a series of tools and workflows that allow to test the functional parts of new replicator sites.

To deploy the FLAME platform in a city, a specific workflow has to be followed and specific tools are to be used. Section 3 introduces the deployment workflow and the ARDENT tool that can be used to test the readiness of an infrastructure, a tool which has been consistently evolving and improving during the project.

During the planning, the validation, and the operation of the infrastructures that host FLAME the replicators have encountered site-specific challenges that required the development of specific solutions. Section 4 exposes these insights and lessons learned from the two initial deployments, as it is likely that similar challenges would have to be addressed in other replicator sites.

Whereas Sections 2 to 4 focus on how FLAME can be successfully replicated, Section 5 deals with another very important aspect which has been key for the successful execution of the open call experiments: FLAME has developed a series of tools and environments in which experimenters will develop, test and improve their services before actually going to a replicator site and running a trial. The established workflow to be followed by an experimenter helps to improve the design of their experiments and allows to detect flaws in said experiment long before actual on-street trials are performed, which reduces the work overhead both for experimenters, platform providers and infrastructure operators.

Finally, in Section 6 this document reports on the replicator progress of the 4 cities in which FLAME is deployed or is to be deployed, followed by the conclusions.

Co-funded by the Horizon 2020
Framework Programme of the European Union

2 ENABLING AN INFRASTRUCTURE FOR FLAME

This section presents the requirements for deploying FLAME in an infrastructure and provides guidelines and suggestions for any new replicator. The methods and tools developed during the project to validate an infrastructure and to set up a FLAME platform are described, as well as how the FLAME requirements have been met in the in the replicator sites at Bristol and Barcelona. Finally, this section describes how to validate the operational readiness of a replicator site.

2.1 MAIN REQUIREMENTS FOR REPLICATORS

There are a series of requirements that need to be fulfilled by replicators that want to deploy FLAME in their infrastructure. This section details these requirements that can be both software and hardware requirements.

2.1.1 SDN Switching Fabric

The switching fabric among compute nodes must be Software Defined Networking (SDN) (OpenFlow 1.3 and above) enabled allowing the FLAME platform to insert the path-based forwarding rules which enables the service routing features described in WP3 architecture deliverables [FLAME-D3.10, FLAME-D3.11]. The switching capability can be realised as a pure software switch using the Open vSwitch[1] implementation running on commercial off-the-shelf (COTS) hardware with any modern Linux-based operating system installed. Alternatively, there are several vendors offering SDN-enabled hardware switches with both physical and optical ports. Across the sites that have the FLAME platform deployed the three switch vendors are in use:

- Pica8 (production discontinued)
- EdgeCore
- Corsa

Across all three vendors PicOS has been used as the operating system which is a Linux-based OS that implements Open vSwitch but translates the switching instructions into the switch's internal hardware (Ternary Content-Addressable Memory (TCAM)) table.

As for the software-based SDN solution, it has been tested on nodes with a 1G network interface controller (NIC) which did not show any performance impact as long as it can be guaranteed the Open vSwitch-based kernel (or VM that hosts the switch) has a single CPU for itself. Tests on an APU2 embedded node (AMD quad-core at 1GHz) demonstrated performance degradations when running Open vSwitch.

The hardware switch has the benefit of guaranteed speeds above 1G (if a 10 or 100G switch was purchased) and the ability to include optical links. However, it must be noted that across all deployments (except Pica8) the hardware switches were configured in hybrid mode allowing standard L2/2.5 switching (MPLS/VLAN/VXLAN/Q-in-Q) as well as SDN-enabled switching. This however lowers

[1] https://www.openvswitch.org

Co-funded by the Horizon 2020 Framework Programme of the European Union

the already size-constrained TCAM table of the switch which holds the rules. And while the path-based rules grow with a constant factor of 2 for each port added to the FLAME Open vSwitch bridge for a software switch, TCAMs have been designed for IP traffic and acts on longest prefixes only, which results in an exponential grow of rules for hardware switches. This results in a physical limitation of how many ports can be added to a hardware switch for the FLAME platform.

2.1.2 Compute nodes

While the FLAME software runs on any COTS hardware that supports Linux some cornerstones can be outlined deciding on the best hardware configuration when acquiring new equipment. For the service routing implementation (FLIPS) the CPU and RAM speed is the key for a performant platform that runs at (1 Gbps) line speed. As FLIPS processes packets in user space the CPU frequency (base, not turbo) is more important than its cache size. A CPU above 3GHz allows FLIPS to operate well. Even though hyperthreading does take a hit on the overall CPU performance FLIPS has been successfully run on compute nodes with hyperthreading enabled.

Another important aspect is to ensure the compute node's vendor has not built in some sort of CPU load-balancing middle-layer which allows an over-provisioning of the CPU threads without a single processing allocating a CPU for itself for most of the time. All FLIPS components have been configured as real-time processes inside the OpenStack virtual machine to ensure they get a fair share of the CPU. When an aforementioned middleware has been installed, processes configured in such a way are hit significantly by a poor performance.

Apart from that, it is also important to choose compute hardware that allows out-of-band control management, often referred to as *lights-out management*. This technology allows a system operator to access the devices remotely, even if the machine is not switched on or if it is unresponsive via *ssh*. Many compute nodes support such features naturally (e.g. the Barcelona nodes supporting the IPMI protocol) and those that do not can be upgraded later on with additional hardware that implements this feature.

2.1.3 Wi-Fi

The radio equipment plays a fundamental role in FLAME, as it is the entry point for users to connect to the media services running in the FLAME platform. There are many radio solutions, not only in terms of different technology (e.g. Wi-Fi at 2.4 GHz, at 5 GHz or even 60 GHz millimetre wave links), but also in types of solutions that can be either commercial or custom.

The radio access network (RAN) solution deployed in Barcelona has been designed from scratch for FLAME. From the beginning, the goal was to assemble hardware with the necessary capacities and performance to serve the FLAME project. This resulted in the construction of a non-commercial, custom Wi-Fi node that includes not only hardware elements carefully chosen by i2CAT, but also a specially configured software stack. The price of each of the Wi-Fi nodes amounts to approximately 1000 €, including the main board, Wi-Fi interfaces, the casing (that includes battery module, alarm module, power over Ethernet, etc.) and the antennas plus all necessary cables.

The specifications of the Wi-Fi nodes are given in D5.1 [FLAME-D5.1]. Here we would like to highlight several key points that are relevant for the replication process. The nodes are equipped with one IEEE 802.11ac compatible Wi-Fi card that supports 2x2 Multiple-Input Multiple-Output (MIMO). This allows for theoretical throughputs of around 700 Mbps, however, in practice we measured throughputs of up to 250 Mbps when using 80 MHz channels. For the FLAME deployment, however, we eventually chose 40 MHz channels, as the unlicensed band in the lower 5 GHz band is quite crowded and from experience, Wi-Fi operating in 80 MHz channels can suffer severe performance issues under such

Co-funded by the Horizon 2020
Framework Programme of the European Union

conditions. Further, the transmission power is always configured to the lower limit (0 dBm), unless the experimenters need additional transmission range for their experiments. Note that the physical configuration of the Wi-Fi nodes is done manually and is adapted for each experiment: experimenters can submit their desired configuration to the Wi-Fi team, which will check whether that configuration can be accordingly adapted.

Each node is running free software, based on Ubuntu 14.04 and supporting all software packages necessary to integrate with FLAME. The key software components are:

- Hostapd: used to manage the Wi-Fi access points (APs)
- Support for VLANs (kernel module)
- Open vSwitch: used to hook up the Wi-Fi APs to the FLAME VLANs and allow for traffic steering with FLIPS

The operation and maintenance of these nodes is completely managed by i2CAT. In case of an incidence, i.e. a Wi-Fi node stops responding, the first step is to perform remote software checks as long as the node can be reached. If not, in the second step the alarm module integrated in the casing of the Wi-Fi nodes is used to force a hardware reset of the Wi-Fi node. If after this reset the node is not yet reachable, the next step is to notify the city council that an intervention on the street is necessary: after a check of the power supply and the network connections (fibre) it is determined whether a crane is necessary to access the Wi-Fi node or whether an issue can be resolved on-ground, e.g. exchanging a cable.

University of Bristol (UoB) has deployed 6 Ruckus T710 Wi-Fi APs in the Millennium Square. The transport and radio parameters are managed manually using the remote Ruckus Software controller. The highest measured throughput has been 116 Mbps on the 40 MHz bandwidth. As the neighbouring cafes and restaurants also transmit their own Service Set Identifiers (SSIDs) on all channels, the FLAME throughput differs across this area. A heatmap is being prepared using a home-made software tool, to identify the most optimum locations and tracks for FLAME experimentation. The towers transmit their unique *FLAMEx* – where x = 1 to 6 – as well as a common 'FLAME' SSID. Depending on the use case, the experimenter can ask for a certain number of SSIDs or certain channels, and request others to be switched off. In most cases FLAME has been carried on 5 GHz frequency, although 2.4 GHz has also been used to optimise performance of a use case. Each AP is configured on the controller via L2 to a flame compute node. Therefore, each FLAME edge service is identified by a VLAN id. For testing purposes, the VLAN can be routed to other APs, e.g. a test AP in the University testbed. The APs receive their IP address via a Dynamic Host Configuration Protocol (DHCP) Server that stores their media access control (MAC) address.

To minimise interruption of this Wi-Fi service, the university of Bristol has also bought a maintenance service contract, that allows real-time communication with technical support in the event of an issue.

The infrastructure in Buseto Palizzolo runs together with the commercial FWA infrastructure from Level7. It must be noted that Level7 is upgrading many of the current links with dark fiber and therefore more channels will be available for the FLAME experiments that will benefit from lower interference issues. The Wi-Fi infrastructure is implemented in outdoor and indoor sites and it is based on Mikrotik devices (both indoor and outdoor) providing coverage across multiple locations. For the indoor locations the Mikrotik RB4011iGS+RM has been used, providing access to both 2.4 GHz and 5 GHz frequencies while for the outdoor devices only 5 GHz frequencies are available (Mikrotik RB921GS-5HPacD, 5 GHz IEEE 802.11ac dual-chain). The devices are advertising a flame specific SSID that is open to experimenters.

Co-funded by the Horizon 2020
Framework Programme of the European Union

It also must be noted that other devices are working on the same unlicensed frequencies. To reduce possible interference, 20 MHz channels are used outdoors, and 40 MHz or 80 MHz can be used indoors). In case a specific experiment needs a "cleaner" coverage in a specific area, Level7 can support this request by providing a different positioning or a narrower antenna (e.g. 30 degrees) in order to get better SNR.

King's College London (KCL) provides an indoor RAN setup where four CISCO Aeronet 3600 Series Wi-Fi access points are deployed in a floor close to a students' area and a lecture hall. The APs are located within the range of each other mainly to demonstrate the effects of handover. The APs broadcast SSIDs with an identifier of the room/office next to it, i.e., *flame-{location}.* Note that there is no DHCP enabled in OpenStack for IP endpoints other than VNFs and the clients must set their own static IP or the VNF provided to supply IP to clients. Again, in most of the cases the APs are operated at 5 GHz however 2.4 GHz can also be used by configuring the controller provided by CISCO.

2.1.4 OpenStack

Across all sites OpenStack is being used to deploy the FLAME platform into the infrastructure in a programmable and automated fashion. OpenStack Ocata has been used in Barcelona and Bristol with OpenStack Pike installed at King's College London, one of the two open call two replicators. The reason of using OpenStack over Open Source MANO or other private cloud solutions such as Kubernetes or Docker Swarm is the wide adoption of OpenStack as the ETSI networks function virtualization (NFV) implementation for telecommunication providers. Also, FLAME demands to specify where each virtual network function (VNF) is deployed, as the entire routing of packets is lifted up to the platform and not done in OpenStack. And that is where OSM failed. The other mentioned solutions above are more targeted at deploying OpenStack itself in an automated manner and do not offer the flexibility and features require by the FLAME platform to be deployed, as they are cloud solutions focusing on vertical scaling of service instances.

As for the compute resource requirements for OpenStack to operate in a production environment, the following table summarises what should be required for OpenStack to operate smoothly.

Table 1: OpenStack compute requirements[2]

Node Type	CPU Threads (vCPUs)	RAM [GB]	Disk [GB]
Controller	8	8	60, HDD
Compute Node	1	1	---

In a production environment it is also recommended to have a dedicated storage host for images. Also, equipping more than one NIC per compute node allows to improve the flexibility in creating networks of various types to support the needs of multiple tenants.

[2] Taken from https://docs.openstack.org/openstack-ansible/latest/user/test/example.html

Co-funded by the Horizon 2020 Framework Programme of the European Union

2.2 EXPERIMENTATION SITE ANALYSIS

The choice of an adequate location for a FLAME deployment is crucial, not only to assure the requirements of FLAME can be met, but also to maximize the impact the FLAME services can have later on during operation for the users and to avoid complications when executing trials. General concepts, such as whether FLAME is deployed indoors or outdoors or whether there are dedicated and suited spaces to host compute and radio equipment, should be taken into account when planning a deployment. This section gives the guidelines for choosing FLAME deployment locations and lists what should be checked beforehand.

2.2.1 Network and deployment planning in the public space

The definition of the network topology and the deployment of both the compute and networking nodes in the street also require careful planning beforehand.

The topology of the network is, among other considerations, closely related to how the experimenter is expected to access the APs in the street: e.g., a centralised access following a backhaul approach compared to a or daisy-chain approach. Whereas a centralised topology forces clients to access the network from a specific AP and also the traversal of packets through the network in a very specific manner until these reach the cabinet's networking and compute nodes, the latter gives the liberty of connecting from every AP and minimises the paths and time taken by the packets to get from the AP to the cabinet or edge node.

As per the planning of the deployment of the nodes to be placed in the street, several considerations are to be taken into account. For one, the Wi-Fi signal attenuation is impacted by many factors: from the expected or usual elements in the street (like trees or the climate conditions, like rain – the latter being covered later) as well as the type of buildings, its materials, the shape of the street and many more. Legal and municipality restrictions have to be taken into account as well prior to the deployment.

Regarding the deployment of the nodes and the links interconnecting the street nodes to that in the data centre, the maintenance work on any of them have to be considered. During the initial definition stage there should be defined a procedure to cope with physical failures in nodes deployed in the street and how to address them. This may require collaboration between the infrastructure maintainers, the municipality office for the on-site services, the external contractors that move to the street to fix the issue and any other actor that may be needed (for instance to block traffic or secure some area during the maintenance window, etc.). Besides the workflow, a separate budget must be kept to finance the works.

2.2.2 Considerations on external factors

When deploying compute nodes or networking equipment on-street, i.e. in street furniture such as lamp posts or cabinets, it is necessary to assure the chosen hardware is resilient to the climatic conditions of the location. High temperatures during summer, high humidity, large temperature changes from night to day, as well as rain are some of the conditions that need to be considered. To give an example, in Barcelona, the on-street compute node deployed in a cabinet was exposed to high temperatures in summer: around mid-day the inside of the cabinet could reach above 60°C, which is

Co-funded by the Horizon 2020
Framework Programme of the European Union

not a critical limit at which the hardware stops working[3], but where it can negatively impact the lifespan of the hardware. Solutions to the temperature issue can be any or a combination of:

- Picking compute nodes of industrial standard that often have a higher resilience to high temperatures
- Installing a cooling system in the cabinet
- Limit operation of hardware inside a cabinet during heatwaves, i.e. turning off any hardware that is not critical for the operation.

Other weather conditions can also easily impact the operability or lifespan of the hardware: the cabinets need to be rainproof and harsh temperature changes should be avoided to prevent condensation effects. As such, it can be critical to make an accurate plan when choosing the location for the compute nodes (and any other on-street equipment).

2.2.3 Indoor vs outdoor deployments

In the FLAME deployments, and following any other 5G deployment, there are two differentiated zones: the core (data centre, indoor) and the edge (APs and cabinet, outdoor). These deployments are interconnected via a high-speed connection (e.g., fibre).

The maintenance of the edge and the links interconnecting the edge with the core are part of the edge deployment and covered in the points above. On the other hand, the maintenance of the computing and networking nodes in the core are planned separately, with no interruption of the public space and considerably less cost associated.

For experimentation, in spite of FLAME trials generally being conducted in the open air, sometimes there can be the need to extend the network to an indoor region, e.g. due to bad weather or events of interest happening indoors. Therefore, as done in Bristol, the network can be extended to an indoor environment that is connected to the FLAME infrastructure. In Bristol, this corresponds to the building next to the Square well known as *We the Curious*[4] (WTC) museum. In such a setup, FLAME SSIDs can therefore be available indoors, using dedicated Wi-Fi APs. In Bristol, there are even other dedicated, smaller indoor spaces where FLAME can be enabled for specific trials, as it is done in the Bristol VR lab, and the Watershed building.

2.3 INFRASTRUCTURE CONFIGURATION

FLAME is quite flexible when it comes to the dimensioning and layout of the infrastructures in which it can be deployed. It is a scalable solution that is not restricted to operating on top of a specific infrastructure in terms of hierarchy (e.g. number and layers of main DCs or edge nodes) or hardware architecture. In order to deploy FLAME in an infrastructure, there need to be certain key enablers available though, such as the capability of the infrastructure to create a slice for FLAME and to assign

[3] The compute hardware tolerates up to 70°-80°C at most. Different hardware may have other limits.

[4] https://www.wethecurious.org/

Co-funded by the Horizon 2020 Framework Programme of the European Union

this slice compute, networking, and radio resources. In this subsection we give an overview of how the different infrastructures (cities) are configured to support FLAME.

2.3.1 Bristol Configuration

In Bristol, the Millennium Square is a public space of about 50 x 50 square meters within Bristol city centre, and a popular visiting place for children and adults. It is also home to several important public events, making it a strong candidate as the FLAME "open testbed". This area is surrounded by ten towers to circulate air in the underneath parking, but also perfect spots to validate wireless technologies. Six antennas tilted towards the square have been set up. Considering wave reflection and interference, not all areas on the square receive the same signal strength.

UoB has deployed the FLAME infrastructure at several separate entities (Figure 1):

1. The data centre resides in the UoB Smart Internet Lab.
2. The Compute Nodes and the OpenStack Controller is placed at the WTC server room, attached to the Millennium Square.
3. Six Wi-Fi APs are on the top of ventilation towers in the Millennium Square, the main venue for events and trials.
4. As a second trial environment for trials, two Wi-Fi APs are installed on the rooftop of the M-Shed museum, radiating at a pedestrian pathway, for backup testing purposes.
5. As a third trial environment, Wi-Fi access networks are deployed inside the WTC network for indoor testing purposes.

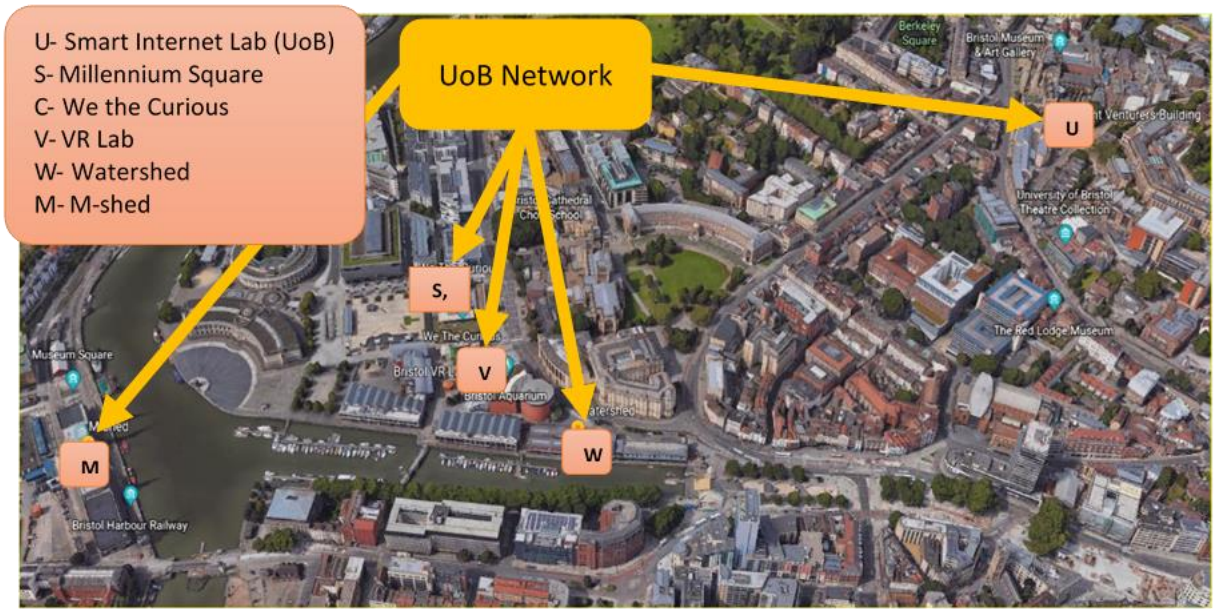

Figure 1: Fibre links deployed in UoB.

The Edgecore switches connect all the compute nodes from each site, UoB lab and Millennium Square. UoB has separated the control plane and data plane from its infrastructure, which means there are two 10 Gbps links, respectively, isolating the control and data planes. These links provide connectivity between the UoB Lab and Millennium Square. While the access to the data centre and WTC sever room

Co-funded by the Horizon 2020 Framework Programme of the European Union

is straightforward, to go through the Wi-Fi APs cherry pickers are needed upon the towers due to the height and the according safety procedures. Below, we list additional information about each element inside the UoB deployment:

- L2/SDN switches: Edgecore running PicOS
    - Two Edgecore as4610_54p and as5712_54x switches working as normal switch with VLAN, trunk and STP capabilities enabled
    - Two Edgecore as4610_54p and as5712_54x switches working as an SDN switch with OpenFlow 1.3 capabilities and IPv6 Label TCAM supporting enabled.
- Seven compute nodes, one OpenStack Controller: Dell PowerEdge R430, 20 CPU cores, 32 GB RAM, 1 TB of disk.
- Ten Wi-Fi APs with 2.4 GHz and 5 GHz frequencies and fast handover IEEE 802.11r capabilities enabled: T710 outdoors; R720 indoor.

Figure 2 shows the components and the connectivity for the OpenStack self-service network installation and one untagged (flat) provider network. In this particular case, the instance resides on the same compute node as the DHCP agent for the network. If the DHCP agent resides on another compute node, the latter only contains a DHCP namespace and with a port on the Open vSwitch integration bridge. UoB fibre network.

Co-funded by the Horizon 2020
Framework Programme of the European Union

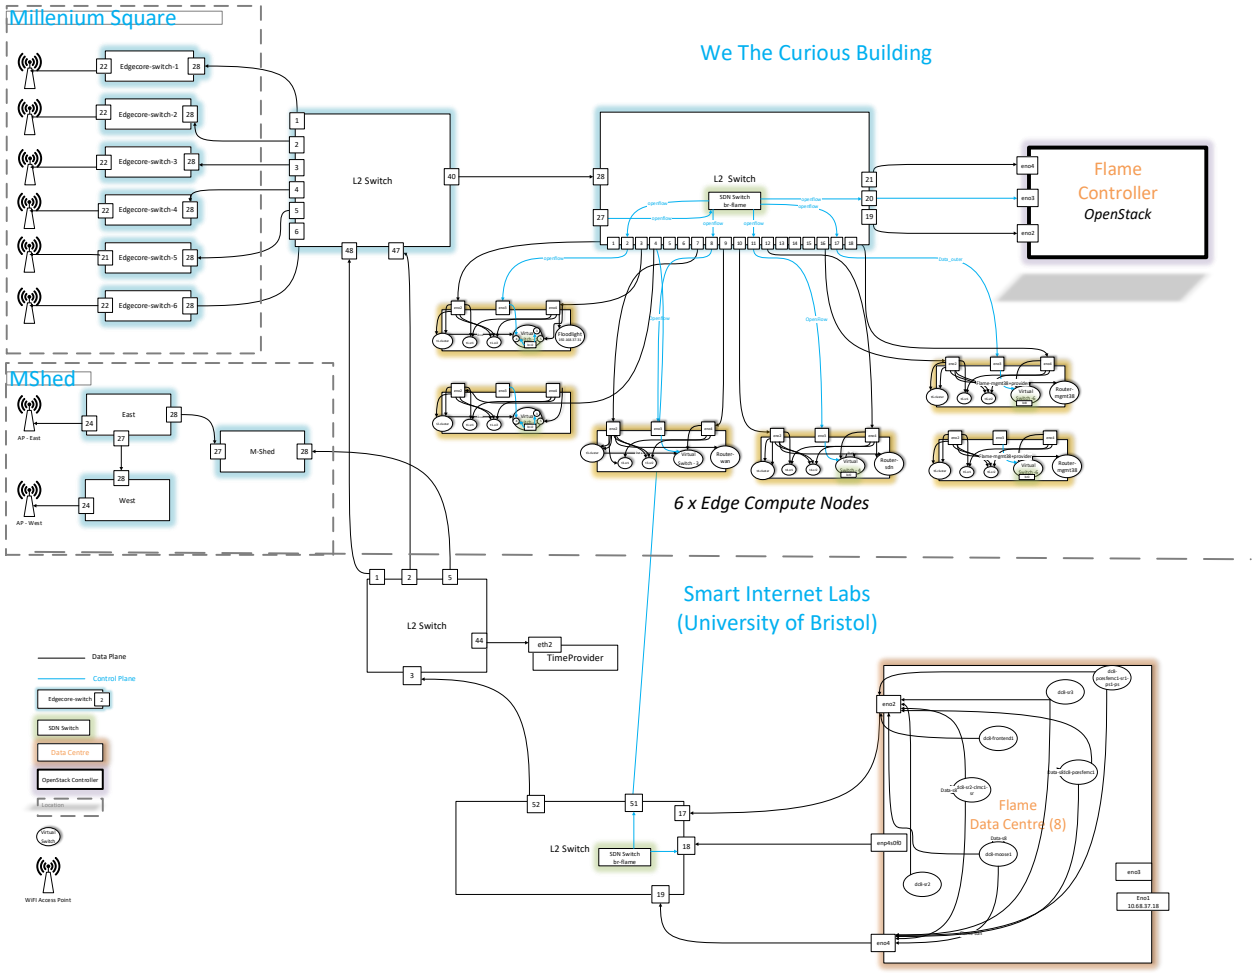

Figure 2: Bristol FLAME topology.

The network was designed according to the requirements of the FLAME platform to provide all necessary internal networks and the user data plane.

Further, to assure high throughput and short delays in the user experience, fibre connectivity between the radio equipment and the compute nodes is recommended. In the UoB setup, each MEC AP is fibred to an individual compute node, whereas in Barcelona each lamp post connects to the same edge compute over dedicated fibre. This ensures FLAME components can be stored right at the edge for optimum performance, and shortest delay.

2.3.1.1 OpenStack Configuration

OpenStack Ocata 3.8.1 has been deployed across the compute nodes, data centre and the controller. It has been configured with the following settings:

- OpenStack Services are installed:
    - Keystone, Glance, Neutron, Horizon, Nova and HEAT.
- All the compute nodes provide 5 wired interfaces: one management interface and four providers interfaces. As a bridging configuration, UoB has used the standard Linux Bridge tool (*/etc/neutron/plugins/ml2/linuxbridge_agent.ini*).

Co-funded by the Horizon 2020
Framework Programme of the European Union

- All the provider networks have been mapped to wired interfaces as follows:
    - Provider 1-> eno2
    - Provider 2 -> eno3
    - Provider 3 -> eno4
    - Provider4 -> enp4s0f0
- As a requirement for the FLAME platform, the OpenStack Installation has been modified to support different MAC and IP address coming from the VMs. The following configuration is necessary as follows:
    - Network Port Security = False

2.3.1.2 EdgeCore Switch Configuration

The FLAME platform operates on SDN switches across the network. At the UoB network, the Edgecore switches have been chosen to handle traffic using normal, as well as OpenFlow capabilities. To accomplish the FLAME platform requirements, several crucial settings have been set on the Edgecore:

1. Pre-define a range of VLAN available to be used by control plane networks
2. Provide OpenFlow 1.3 with IPv6 full capabilities (IPv6 Label)
3. Increase the size of the TCAM memory available to be used strictly by the OpenFlow rules

The FLAME platform is a distributed network, once the deployable computing resources can be placed either in the edge or the core. In the OpenStack controller, FLAME networks along with VLANs and interfaces are set up according to specific FLAME requirements. The HEAT template creates the hierarchical topology, where the service routers (SRs) and the clusters as part of the edge computing. There is a switch between one compute node and one tower that is providing the Wi-Fi access. The HPN Lab data centre also has SRs and Clusters, but additionally it has been used to place the CLMC, SFEMC, PS and PCE components (details in the replication progress, Section 6.2.2).

2.3.1.3 The Floodlight Controller

The Floodlight controller managing the SDN fabric has been set up in the FLAME OpenStack controller. In the figure below the controller has created a topology of the SDN fabric, and the SDN switches running on Edgecore switches.

Co-funded by the Horizon 2020
Framework Programme of the European Union

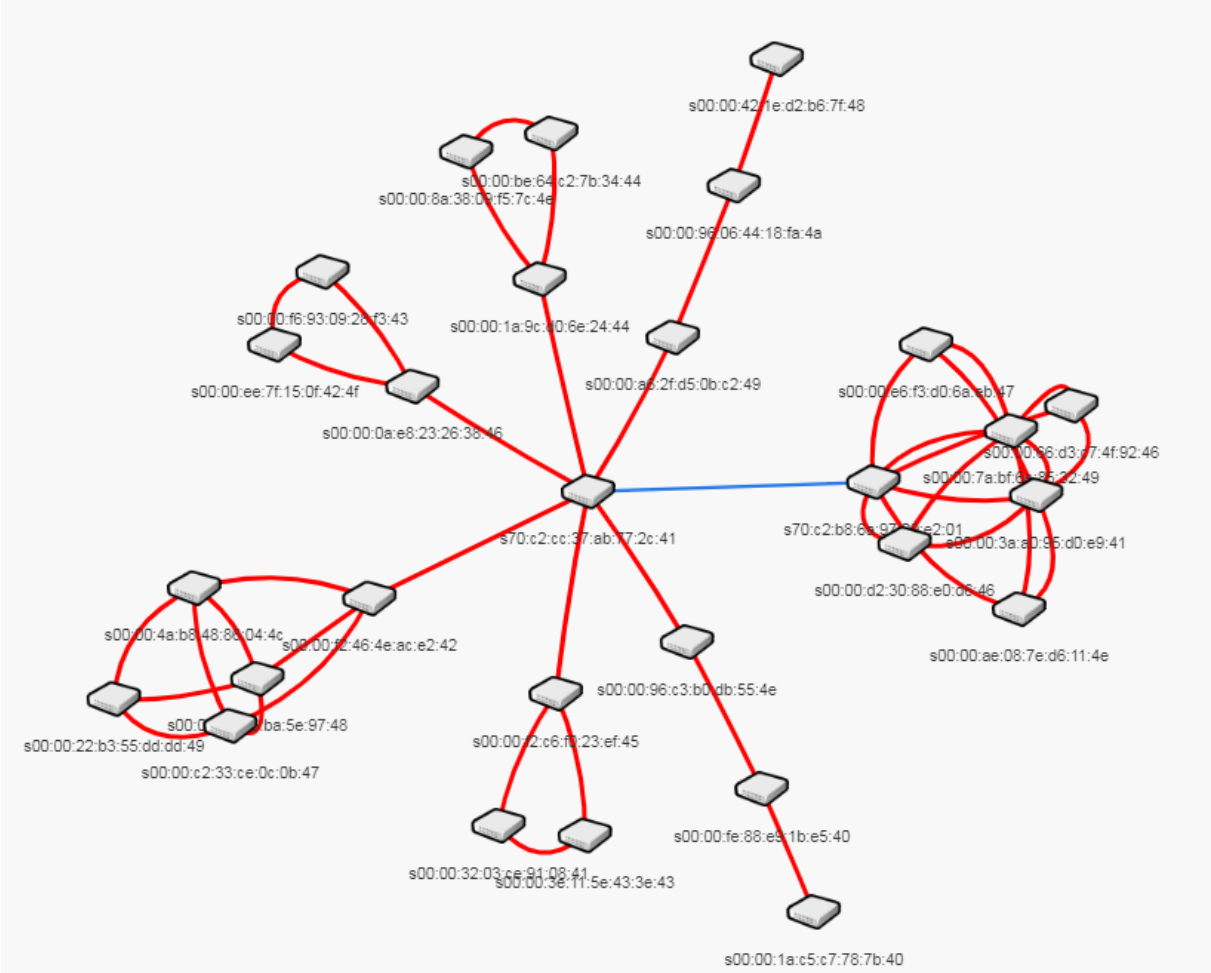

Figure 3: Bristol SDN topology.

The topology view accessible from Floodlight is a very useful tool to monitor the SDN fabric and locate possible broken links. Figure 3 depicts the SDN topology when the FLAME platform is fully operational in the UoB infrastructure.

2.3.2 Barcelona Configuration

The Barcelona FLAME deployment can be divided into two main sites, one offering an edge compute node as well as RAN capacities and one data centre (main DC) site. The two sites contain the servers forming the cluster and the switches interconnecting them inside the data centre or across the city. In Barcelona there is one cabinet (edge) with the fog server and a router and one data centre (main DC) with a cluster of three servers and, connecting both, dark fibre deployed throughout the city to connect both sites. The Wi-Fi nodes deployed at the edge connect to the edge cabinet over fibre.

The compute nodes and Wi-Fi devices are connected over a L2 network operating on top of fibre and copper, using a range of VLANs in a VLAN trunk that is configured in the switches controlled by i2CAT and placed across cabinets and i2CAT data centre, as well as the edge. This connects the servers (and thus any virtual instance inside them) at layer 2, disregarding the physical location (cabinet or data centre) and allows traffic towards or from the UEs to reach services running in the edge or main DC.

The slice configuration in which FLAME would operate was elaborated as a common effort between i2CAT and Interdigital, first identifying all platform and infrastructure elements that would be

Co-funded by the Horizon 2020
Framework Programme of the European Union

deployed, followed by a planning of the VLANs to enable aforementioned L2 network connecting these elements. As an update from the configuration showed in previous deliverables, the final logical design of the FLAME topology assigns dedicated clusters to each Wi-Fi node and also integrates emulated user equipment (see Section 6.3.2 for details).

The resulting VLAN assignment for Barcelona is depicted in Figure 4, with assignments reaching from the Wi-Fi nodes (*Gateworks*), over the Cabinet server, to the main DC (*Omega*). It is noteworthy highlighting that in Barcelona an additional degree of complexity was introduced, as there is only a single edge server, but with the help of adequate VLAN assignment, each of the lamp posts was logically connected to a dedicated cluster each. This required an additional engineering effort that allows platform users to instantiate localized services for each lamp post. Further, for the Wi-Fi APs a custom solution is used. The single board computers equipped with Wi-Fi interfaces running on Ubuntu 14.04 (Gateworks) have been designed from a hardware and software point of view to fully integrate with the FLAME deployment. The key element to enable the Wi-Fi APs to integrate is the use of virtual switches (Open vSwitch) to extend the SDN fabric from the wired infrastructure to the lamp posts and thus the Wi-Fi APs.

Please note that to each VLAN displayed in Figure 4 the value “1500” needs to be added to get the actual VLANs used in the Barcelona deployment (e.g., in SR-1 for Gateworks 0, the first three links – access, data and SDN- are set to “10”, “11” and “70”; which shifted to the actual VLANs would be “1510”, “1511” and “1570”). As such, for the connectivity to the lamp posts, to connect the data clusters elements and to enable the user data to traverse the infrastructure, OpenStack communicates with the networking by defining N (where N is the number of APs) access networks and N+2 of data networks, each of them related to a VLAN:

- Access networks: management addresses to reach platform nodes. VLAN: 1510, 1520, 1530, 1540
- Data networks: data plane for platform nodes. VLANs: 1511, 1521, 1531, 1541, 1551, 1600
- SDN control: where platform SRs are interconnected. VLAN: 1570

Co-funded by the Horizon 2020
Framework Programme of the European Union

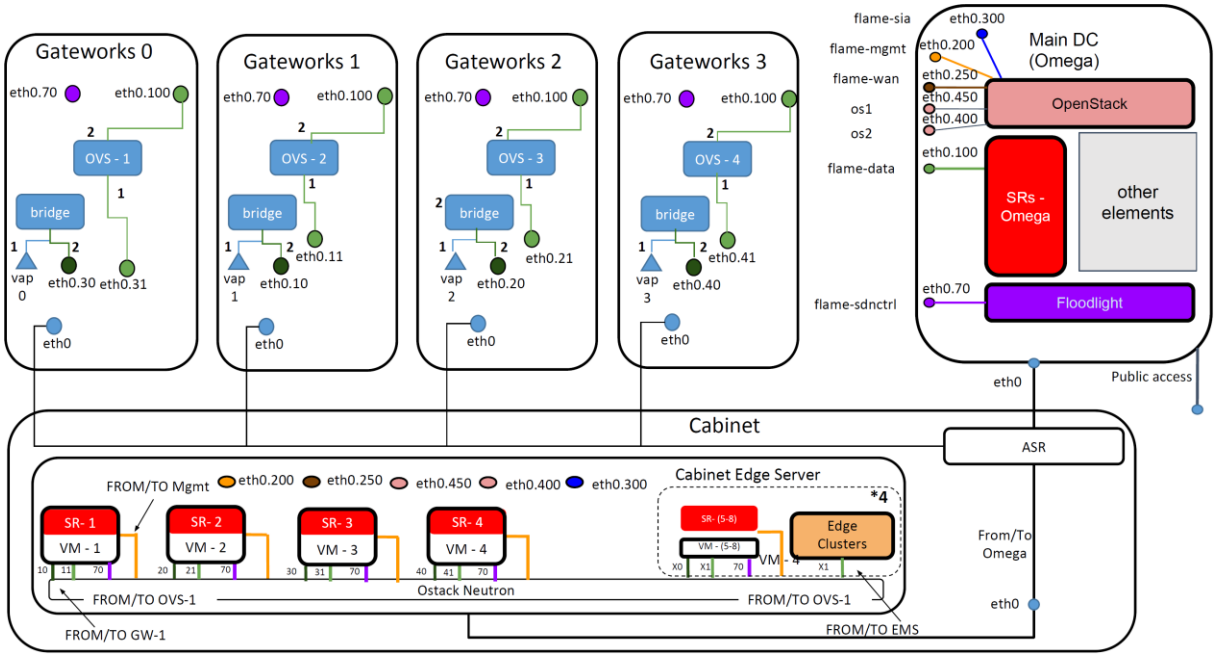

Figure 4: Final VLAN configuration for the Barcelona infrastructure.

The virtualization stack that takes care of the VLAN configuration of the edge and compute elements is provided by OpenStack Ocata. In the case of Barcelona, there are as many availability zones (AZ) as locations: one for the core, with the main DC compute nodes (i2CAT Omega Building) and the one at the edge, collocated with the RAN equipment. The HEAT templates used during the deployment of the platform select the proper AZ to instantiate each node. Figure 5 shows the different areas (core to the left, edge to the right) and how these are connected across the city through a hierarchical set-up of switches and routers.

The FLAME platform is deployed in one of the computing nodes of the OpenStack cluster. A “seed” VM is reserved and inside it. HEAT templates are used to instantiate it via some appropriate scripts. The platform itself pre-allocates resources from the system (CPUs, virtual memory, disk) so that the experimenter can deploy an SFC through the portal provided by the platform.

Co-funded by the Horizon 2020
Framework Programme of the European Union

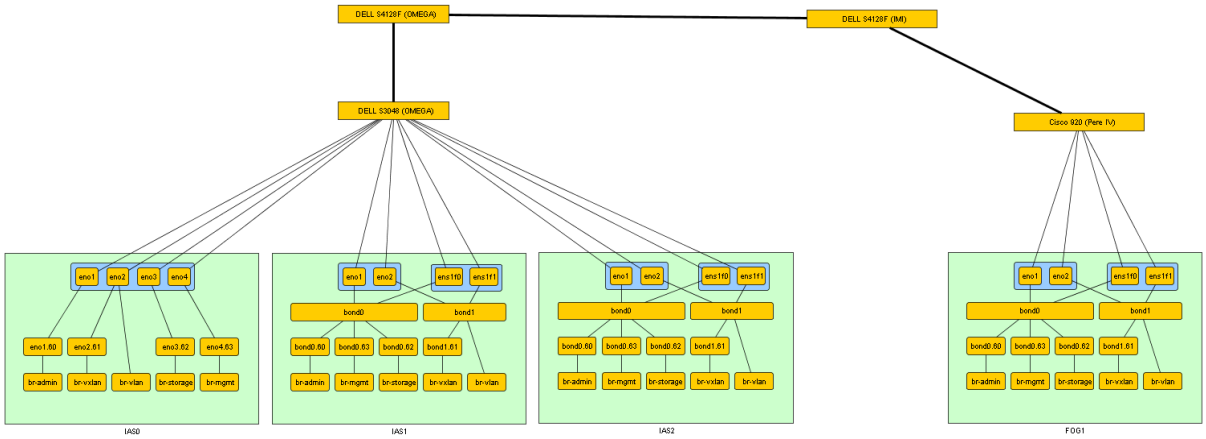

Figure 5: Network topology and availability zones for the computing cluster.

As per the Wi-Fi nodes in Barcelona, a proprietary solution is used; where RAN connectivity is provided by running hostapd on top of the physical Wi-Fi interfaces. To integrate these Wi-Fi nodes with the FLAME platform, Open vSwitch bridges are used to attach the nodes to the required access, data and SDN control networks. Figure 6 depicts the solution applied for Gateworks 0, which can be applied to any of the Wi-Fi nodes.

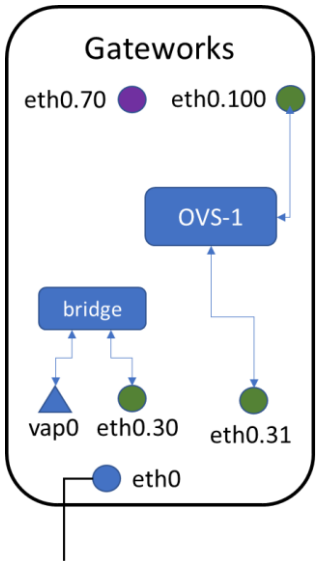

Figure 6: Schematic of the Open vSwitch and VLAN configuration of a Barcelona Wi-Fi node.

A dedicated hostapd process manages the interface named *vap0*, a virtual AP that is generated on top of the physical Wi-Fi transceiver, which is configured to operate in the lower 5 GHz band with a 40 MHz bandwidth (IEEE 802.11ac). In Barcelona 4 nodes are deployed and 2 configurations in the lower 5 GHz band are possible at 40 MHz: using channel 36 as base or channel 44. As such, we chose to assign channel 36 to Gateworks 0 and 2, and channel 44 to Gateworks 1 and 3 to reduce the interference between neighbouring lamp posts.

On each Gateworks, the Open vSwitch bridge named *bridge* connects the virtual APs with the access network. In the example on the left, this corresponds to attaching *vap0* and VLAN 1530 (via eth0.30)

Co-funded by the Horizon 2020
Framework Programme of the European Union

and acts as a self-learning bridge [5]. This allows the exchange of packets between users attached to the AP and the SR element that is attached to the 1530 network and is located in the edge cabinet. There the traffic switches from the IP domain into the FLIPs domain and is sent back to the Wi-Fi node. For that, the second VLAN interface on the node is attached to eth0.31, that leads to the VLAN 1510 data network in the FLIPs domain. The responsible Open vSwitch bridge for handling this traffic and forwarding to the right destination is managed by a Floodlight controller instance (sitting on VLAN 1570, control via eth0.70). The Open vSwitch bridge is completely managed by the Floodlight instance and rules are provided by the path compute element as part of the FLAME platform. From the VLAN 1600 (corresponding to eth0.100), the traffic can then reach the rest of the infrastructure, e.g. other Gateworks nodes, as well as the clusters at the edge or the main DC, where the FLAME services are running.

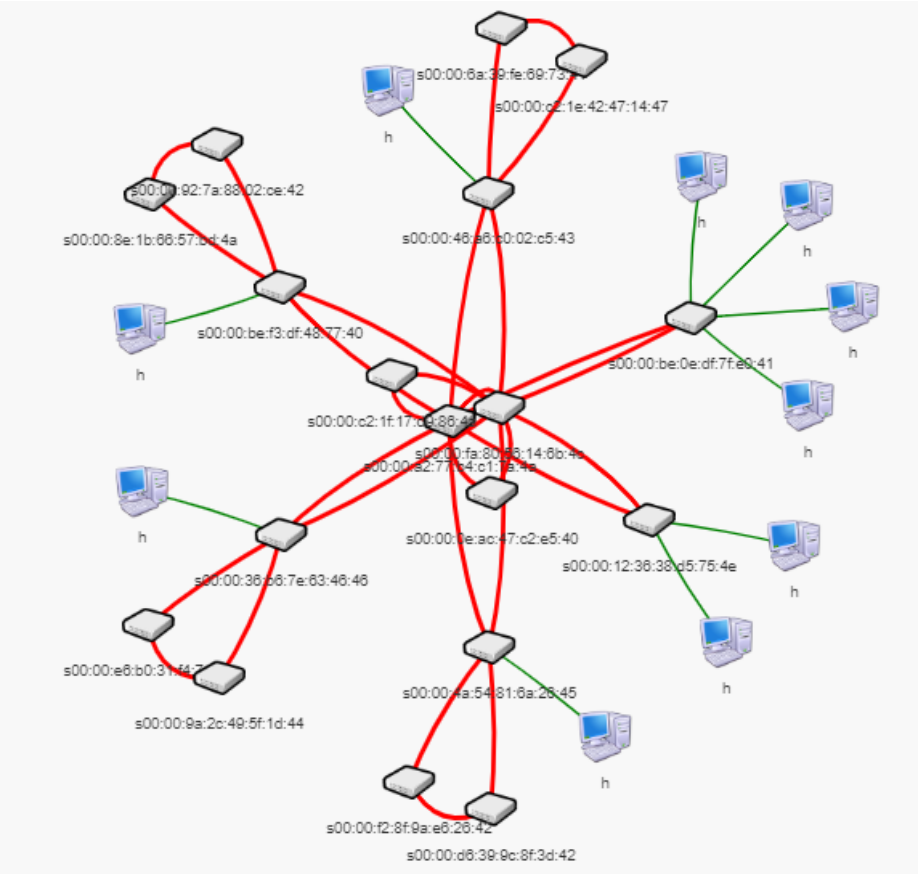

Figure 7: Barcelona SDN topology as seen by the Floodlight SDN controller

The resulting connectivity map of SDN switches and service endpoints during a full platform deployment in Barcelona is presented in Figure 7. This view corresponds to the topology as seen by the SDN controller during an experiment.

2.3.3 Busseto Palizzolo Configuration

The deployment in Sicily has a main site in Buseto Palizzolo (a rural and isolated community) as well as a secondary node in Level7 main offices in Palermo. The node in Palermo has been made available only

[5] As noted before, "1500" needs to be added to the VLANs shown in this figure to obtain the actual VLAN.

Co-funded by the Horizon 2020 Framework Programme of the European Union

for testing purposes and to speed up the implementation of experimenters on Level7 infrastructure, while the real and main infrastructure that will host experiments from third parties is located in Buseto Palizzolo, with indoor and outdoor installations.

It also must be noted that the node in Palermo has been planned and implemented in order to provide technical availability of the Level7 FLAME infrastructure in case that i) the infrastructure in Buseto Palizzolo is not accessible or ii) the experimenter wants a technical test with real devices done by Level7 personnel before going to the real infrastructure with experimenters. It also permits technical tests by Level7 personnel without going to the Buseto Palizzolo site (around 1h is required for the commute to the site from Level7 offices).

The implementation of the OpenStack main infrastructure has been done in Level7 offices in Palermo with 3 dedicated servers. These servers are connected to Buseto Palizzolo via a 10 Mbps link with 50ms round-trip time due to the fact that Palermo and Buseto Palizzolo are connected via a geographical link. This makes the edge computing scenario more realistic and relevant, where the servers implementing the edge computing in Buseto Palizzolo must provide services locally without accessing "central cloud resources" from the main office in Palermo.

In the future, the bandwidth of this link could be upgraded in order to provide a better connectivity between the two sites. However, even if this link will be upgraded to higher bandwidth, due to the network architecture the round-trip time between the two sites will still be around 40-50ms.

2.3.3.1 OpenStack Configuration

Level7 has implemented OpenStack Ocata on 3 dedicated servers with the following configuration:

- 1 controller node
- 2 compute nodes (compute-1 and compute-2)

The Openstack services that have been installed are:

- Keystone, Glance, Cinder, Neutron Horizon, Nova, and HEAT

The Cinder service is actually running only on compute-1 with a dedicated storage of 1000 GB but it can be easily installed on other nodes as well as expanded, if needed.

Description	vCPUs	RAM	Storage
Controller[6]	N.A. (16)	N.A. (71 GB)	N.A. (1.8 TB)
Compute-1	16	141 GB	930 GB (Cinder) + 8 TB (Glance)
Compute-2	16	141 GB	7200 GB (Glance)

[6] The resources are not available to FLAME but are used by Openstack for the controller node

Co-funded by the Horizon 2020 Framework Programme of the European Union

2.3.3.2 The Floodlight Controller

The Floodlight controller managing the SDN switches has been installed, as a separate VMWare machine on another cluster outside the Openstack FLAME hardware, and it is controlling the SDN layer of the Level7 installation and communicating with the OVSwitch instances.

2.3.3.3 Node configuration in Palermo (Level7 Offices)

One FLAME node has been implemented in Palermo, in order to speedup technical tests with real hardware or provide access to Level7 infrastructure, in case i) a specific node of Buseto Palizzolo cannot be accessed or ii) in order to test on real hardware in "the same day" without the need to go to Buseto Palizzolo. Obviously, the tests on this node must be orchestrated together with Level7 personnel who can provide support from remote but it cannot have the same "social" impact or user experience that can be obtained in Buseto Palizzolo due to the lack of real experimenters.

From the technical point of view, the node has the same hardware features that are present in many nodes in the real infrastructure, i.e. a server with many cores and an access point from Mikrotik (i.e. RB4011iGS+RM with 2 GHz and 5 GHz antennas, the same model is used in the indoor environments in Buseto Palizzolo). The compute server is part of the Openstack installation and it has 24 vCPUs, 32 GB RAM, 1.8 TB of storage.

The Wi-Fi AP is directly physically connected to the server. Out of band control (i.e. direct access to the device) of this device is obtained connecting it to a switch that is running the "out of band" network for the FLAME infrastructure, i.e. separate address space that is available only to Level7 and not to the experimenters.

2.3.3.4 Buseto Palizzolo Infrastructure Configuration

The infrastructure in Buseto Palizzolo is located in a rural area that can be reached in 1h driving from Palermo. The infrastructure is made of indoor and outdoor nodes and it is located in public places (school, square, museum, library, etc.) in order to make the experiments as much as effective as possible, with the help of the local community.

Every node is composed of:

- one compute resource (a dedicated server)
- one switch that is used to map the VLAN and to turn on/off the access devices (via PoE)
- one or more access devices that are deployed indoors or outdoors.

The nodes are currently connected via a dedicated radio link that should be upgraded to dark fiber in 2020. One SDN OpenVSwitch (based on x86 hardware – 16 vCPUs, 16 GB RAM, 450 GB SSD) is located at the cemetery in a "star" topology, i.e. all the nodes are (or will be) connected to this node.

The current topology, of the nodes, is illustrated in Figure 8.

Co-funded by the Horizon 2020
Framework Programme of the European Union

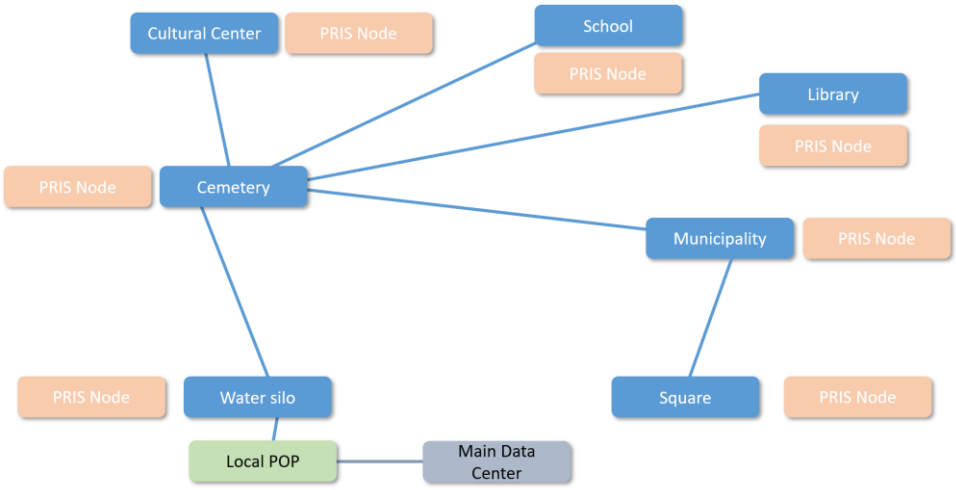

Figure 8 - Buseto Palizzolo current topology

It must be noted that, where it is shown a "PRIS node" it is intended that a compute node will be locally available. In the topology, it is intended that at the cemetery the SDN OpenVSwitch is installed for SDN features. For the computational nodes, most of the nodes are based on Dell hardware with 24 vCPUs, 32GB RAM and 450GB SSD. For some nodes (e.g. Water silo and cemetery) AMD based hardware has been implemented (8 vCPUs, 16 GB RAM, 450 GB SSD).

2.3.4 London Configuration

Within KCL Strand campus, which is conveniently located at the heart of London, the Wi-Fi APs are in the first floor of the building. The APs are deployed closer to a PhD open office, a small lecture hall, staff area and one in the 5G lab (as shown in Figure 9). The floor is mainly accessed by students and staff members of the center for telecommunication research. This makes the King's College testbed a unique candidate to test an indoor set-up of FLAME with applications especially meant for educational media delivery. The other unique feature of the testbed is that there is an already established link via the Slough exchange point to UoB testbed through the 5G exchange located at slough.

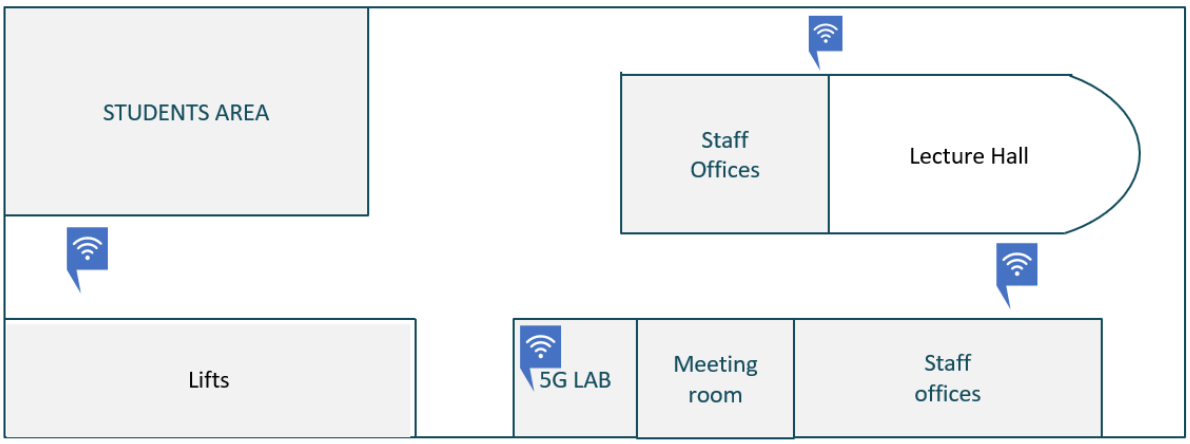

Figure 9: Location of Wi-Fi Access Points deployed at King's College London Strand Campus

KCL FLAME infrastructure set-up is quite similar to that of UoB's with following elements:

Co-funded by the Horizon 2020 Framework Programme of the European Union

1. Four Wi-Fi APs located at King's College London, Strand Camps.
2. The datacenter is located at the lower ground of King's building, where the compute nodes are located.
3. The Openstack controller is also located in the datacenter.

The core-switch connects Wi-Fi Aps to the edge computes located at the datacenter, as depicted in Figure 10. While two of the computes are standard hosts, the other two are setup as real-time hosts in order to test specific applications that require the real-time features. Independent KVM nodes are also made available to test the performance of applications under different scenarios. The individual components made available to deploy FLAME win KCL testbed are as follows:

- An Edgecore AS4610-54T running in Openflow mode is dedicated to the project. The switch features 48 x 1 Gbps ports and two SFP+ ports.

- Four compute nodes with identical hardware configuration. The server model is Dell R630 with 88 vCores, 128 GB RAM, 1 TB storage, SFP+ and Gigabit NICs. The compute nodes are managed by OpenStack.

- Four Wi-Fi APs are CISCO Aeronet 3600 Series, configured with a VLAN and SSID slice dedicated to FLAME.

- The infrastructure carrying FLAME traffic also includes an Edgecore AS4610-54P where the CISCO APs are connected, an Edgecore AS5812-54X which is the 40 Gbit core switch of the 5G testbed and a Corsa DP2100 which is Openstack's main provider switch that connects VNFs to the rest of the infrastructure.

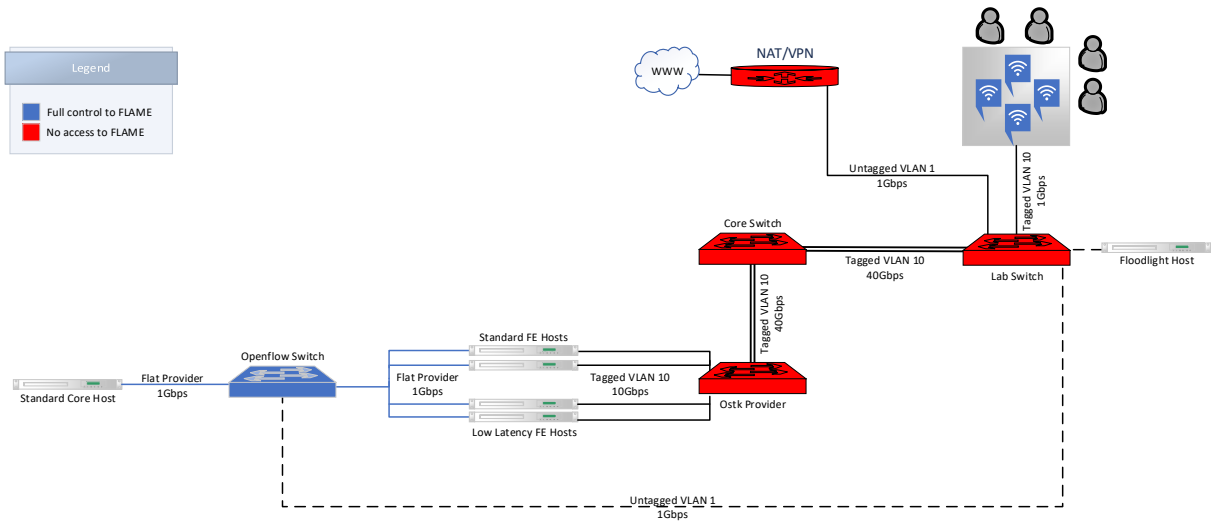

Figure 10: KCL Infrastructure set-up

The networks were setup based on the requirements of FLAME alongside the necessary internal networks.

2.3.4.1 OpenStack Configuration

Openstack Pike is used to manage the compute nodes and has been configured with full redundancy for routing and DHCP services.

Co-funded by the Horizon 2020
Framework Programme of the European Union

Furthermore, the routing configuration is fully distributed which provides increased performance for VMs attached to overlay networks since there is not a single router residing on the controller node to handle all the VM traffic. The controller node can store up to 1TB of VM images, with the average size of a VNF image being around 2 GB. Images are inflated when they are launched and the user can configure the running instances total storage, however, it is recommended to keep it below 10 GB to avoid problems with block allocation during instantiation

Block storage using Cinder is provided using LVM backend. This allows the user to more easily manage VM images and create snapshots and backups. The other services that are installed includes: Keystone, HEAT, Glance, Neutron, NOVA and Horizon.

Four host aggregate groups have been configured in Openstack to accommodate standard VM hosting, real-time VM hosting, standard containers and real-time containers. Each host aggregate group defines "hypervisor" and "real-time" directives to determine where a VM should be created.
On the individual physical hosts the 10GbE interfaces assigned are as follows: 1 – used for Openstack API, 2 – used for provider network, 3 – used for storage network, 4 – used for overlay network.

2.3.4.2 Switches Configuration

The FLAME traffic is carried through three switches to the Wi-Fi AP (as shown in Figure 10) –

(i) The Lab switch (Edgecore 4610-54P) for connecting Wi-Fi APs to the Datacenter switch. Aggregate link *ae1* connects to the Datacentre Core switch at 40 Gbps over 4x10 Gbps Link Aggregate Ports (LAG) at the corresponding *ae7* interface.

(ii) The Core Switch (Edgecore AS5812-54X) which acts as the datacenter core switch. VLAN 10 added to aggregate links ae5 and ae7, where ae5 connects to the Openstack Provider network switch at 40 Gbps over 4x10 Gbps LAG. Aggregate link ae7 connects to the Lab Switch where Wi-Fi APs are attached on the corresponding ae1 interface.

(iii) The Provider Switch (Corsa DP2100) which is an Openstack provider switch and is configured with passthrough mode. The bitrate per interfaces is guaranteed at 3.33 Gbps and the switch allows all VLANs to pass from external networks.

2.4 OPERATIONAL READINESS

In order to validate whether an infrastructure is ready to host FLAME, some tests are necessary to assure that all physical and logical elements are working correctly. This section introduces these procedures and describes how they have been applied to the different replicators.

2.4.1 Slice Creation

The creation of the slice for the FLAME tenant is required on various levels in the infrastructure such as:

- OpenStack: A FLAME project is required with one user. Furthermore, the segmentation identifiers for external provider networks must be aligned with the VLAN configuration of networks outside of OpenStack

Co-funded by the Horizon 2020 Framework Programme of the European Union

- SDN switching fabric: A dedicated bridge with (access) ports must be configured enabling the logical topology envisaged.

- Access networks: These networks allow IP endpoints to access the deployed platform and user traffic must be carried to the correct compute nodes where the platform can handle it. This part of the slice probably requires VLANs as well.

- Radio access: If a wireless access is envisaged (e.g. Wi-Fi) the usage of virtual APs is mandatory to separate the FLAME slice from any other tenant.

The remainder of this section presents the approach of how the slice has been created in Barcelona and Bristol and which tools and/or methodology they have been using to ensure the organisation of shared information across technologies.

2.4.1.1 Barcelona

The Barcelona deployment does not relay on an SDN-enabled networking fabric, but instead it uses a "hybrid deployment"; consisting of a cluster of computing nodes (where the FLAME platform runs) distributed throughout the city and managed by OpenStack. These computing nodes are connected to each other via some legacy, non-SDN switches and routers. To interconnect them, the *cluster has a pool of VLANs* at its disposal *that can be* manually *assigned by the OpenStack operators to any newly created network* during the deployment process and the integration with the FLAME platform. *The pool consists of the VLAN*s 1500 *to* 2000*, both included. Every packet* sent from an instance running in the OpenStack cluster will be tagged appropriately (that is, according to the VLAN defined as "segmentation ID" in each OpenStack virtual network).

As previously shown in *Figure 5*, in the core zone such packets will be received by the Dell S3048 switch; whilst in the edge zone the packets will go through the Cisco ASR920 router. This means that both network devices receive the tagged L2 frames from their closest compute nodes. To be able to receive the frames, identify them from a specific source and maintain different networks throughout the different zones, the network devices and equipment must be previously manually configured to have tagged all the VLANS from the available range in "trunk" mode, for both uplinks and downlinks. Then, each interface has its VLAN(s) assigned as best deemed. Besides that, other configurations and controls are in place (for instance, the Spanning Tree Protocol is enabled in the switches so as to avoid loops). In this way, the different "flame-access", "flame-data" and "sdnctrl" networks connect all computing nodes in a transparent manner.

2.4.1.2 Bristol

The UoB deployment utilizes SDN-enable networking fabric and virtual switches, in a "hybrid deployment." The environment provides a cluster of seven compute nodes managed by OpenStack and distributed across the city as shown in Section 2.3.1. Six compute nodes are placed around the Millennium Square and serving each tower. Each tower has the LTE and Wi-Fi capabilities enabled. One compute node is placed at the UoB HPN lab. The cluster has a pool of VLANs (from 50 to 80) for the control plane, which manually can be assigned by the OpenStack Controller. There is a dedicated data network for all compute nodes. Therefore, every packet using the data plane is sent from an instance running in the OpenStack Cluster and cross fabric and virtual switches. There are two Edgecore switches, one placed at the Millennium Square and other at the HPN Lab.

According to the Bristol FLAME Topology (as shown in Figure 2), in the UoB HPN Lab data centre and the Millenium Square the packets are received by an Edgecore as4610_54p switch, respectively. Both

Co-funded by the Horizon 2020 Framework Programme of the European Union

switches have the OpenFlow 1.3 capabilities enabled and are interconnected by two optical links of 10Gbps (control plane and data plane). For the control plane, each network slice has its VLAN(s) assigned as best deemed. Moreover, the control plane set up is related to the "flame-LAN," "flame-mgmt", and "flame-sia" networks, which transparently connect the computing nodes. Besides, the data plane set up is providing access to the OpenFlow switches, which, on the OpenStack, is given by the "flame-data" and "flame-data-outer."

2.4.1.3 Buseto Palizzolo

The implementation of FLAME in the Buseto Palizzolo testbed follows a hybrid deployment, where parts of the network are SDN based (the SDN OVSwitch at the cemetery) while VLAN with full bridging are used in other parts of the testbed. The OpenStack main deployment is in Palermo, which is connected to Buseto Palizzolo, via a geographical data link. This makes the Buseto Palizzolo a real edge computing scenario, where each single node in Buseto Palizzolo is responsible to provide the services, locally, to the users via the direct connection to the access devices or using neighbour nodes.

In order to deploy FLAME experiments in a separate environment (Level7 is operating a parallel commercial infrastructure in Buseto Palizzolo) the access radio network has been realized with dedicated devices. The fact that each radio device is connected to a computational node, makes the data processing easier in many scenarios.

2.4.1.4 London (KCL)

KCL deployment for FLAME carries four physical hosts managed by OpenStack and the infrastructure utilizes the SDN-fabrics. A dedicated data network is allocated for each of the four access points to the Wi-Fi access points. It is worth noting that KCL infrastructure has multiple tenants in terms of both hardware and software. The FLAME platform is collocated with other tenants and the access is given to certain part of KCL's infrastructure (as described in Section 2.3.4).

The control plane in FLAME, especially for flame-mgmt and flame-sia networks, is configured in such a way that the connection to compute nodes is transparent. For example, VLAN 698 is pointed to ae5 such that flame-sia can directly talk to outside.

2.4.2 Slice Readiness

Once the slice has been created is must be tested against its operational readiness. The configuration of the various infrastructure components deems successful (operational/basic connectivity) when testing if they work using ICMP messages or any other echo request – response protocol. However, it is of most importance to stress test the infrastructure slice over a significant amount of time to demonstrate its readiness. For instance, faulty hardware or software very often reveals its true readiness when put under load. Especially switches, network interface cards, drivers, and VLAN configurations are candidates for a thorough testing beyond simply echo request response packets.

As outlined in detail the ARDENT specification in D3.11 [FLAME-D3.11], the data plane of the platform traverses internal and external OpenStack networks which are partially configured with standard IP routing (e.g. between SRs and clusters) or with the novel service routing which uses path-based forwarding rules (e.g. between SRs). Both types of networks must be stress tested first to ensure the nominal line speeds can be achieved. In order to conduct these tests without the platform being

Co-funded by the Horizon 2020 Framework Programme of the European Union

deployed, it is necessary to fall back to IP-based benchmarking tools such as Iperf[7]. All SDN-enabled bridges that are part of the slice should receive appropriate IP rules to allow packets to traverse the switch.

Across all access and data networks Iperf TCP should be run first for a period of 5min or longer. Assuming an MTU of 1500 resulting in 1518 octet frames should result in

a) approximately 98% of the nominal throughput. For instance, for a 1G link 970 – 980 Mbps should be reported after the contention window has become stable. Using Iperf3 will provide the CWD value as part of the standard output.

b) Stable contention window once is has stopped growing

Then, Iperf UDP should be configured with the average rate reported with TCP and the same throughput numbers with a packet loss of << 0.1% must be observed.

If any of the key performance indicators (KPIs) given above cannot be met a thorough investigation must be conducted and the issue rectified that stops the infrastructure from operating at nominal speed.

2.4.3 SDN Underlay

Given a positive slice readiness testing, as described in the previous section, the SDN underlay must be verified against its operational readiness in relation with the path-based forwarding rules. In order to achieve that an Iperf-like application is available as part of the SFR implementation that allows to stress test connections among SRs. While measuring the throughput per second this application can also measure the round trip time for a configurable amount of packets per second. The results must be identical with the Iperf measurements obtained in the previous section.

2.4.4 Wi-Fi Access Networks

The testing of a wireless link operating in an unlicensed band is a challenging task and this section outlines the methodologies developed in FLAME to achieve that. As for the testing of wired links, the nominal speed must be determined by ensuring the software and hardware used to create the AP as well as the hard- and software used on the client attaching to the AP support the Wi-Fi standard configured. Starting with IEEE 802.11n MIMO radios made their way into the IEEE specification and whether 2x2 or 4x4 (or even higher combinations) of MIMO radio are being used affects the nominal achievable throughput numbers significantly. Also support for channel bandwidth configurations must be checked on both the AP and the client. Bearing these hard numbers in mind, the location of where the client is located (line-of-sight, potential reflections of walls, obstacles like trees with or without leaves, time of the day, weather conditions) and how its antenna had been directed towards the AP will need to be considered.

Taking the points above into account, it is worthwhile analysing the area that is covered by all APs for potential point of interests where users will most likely access the platform. This allows to identify the point where the KPIs throughput and round-trip time should be measured.

[7] https://iperf.fr/

Co-funded by the Horizon 2020
Framework Programme of the European Union

2.4.5 Virtual Networks in OpenStack

Determining the upper limits of performance of VM and compute node configurations in a city can be elemental to understand the performance experienced by users, but also the performance between the different virtualized elements in a FLAME deployment. Since virtual networks are used to connect service VMs or even emulated UEs, testing the limitations of these virtual networks should be tested to determine upper performance limits and also possible configuration issues. The required tests can be executed with basic tools like Iperf that allow to generate UDP and TCP streams between any two endpoints of a FLAME deployment.

The following tests should be performed:

- VM in one cluster against another VM in the same cluster, i.e. on the same compute node. This test allows to determine whether there are any issues in the virtualisation configuration of the VMs that are being deployed. In an optimal case, running throughput tests between two VMs that are deployed in the same physical compute node deliver very high throughputs of up to tens of Gbps. If the measured throughput is clearly below such rates, then a configuration of the VM could be causing trouble, which eventually could reflect in the performance experienced by the users.

- VM from one availability zone (in Barcelona edge or core) against a VM in another availability zone. Testing the communication between different availability zones allows to determine whether the underlying infrastructure that connects the two zones (L2 network) is capable of carrying the nominal load. In Barcelona, a nominal end-to-end connectivity of 1 Gbps is expected[8], thus any Iperf tests executed between two VMs needs to reach throughputs of above 800 Mbps.

- Wi-Fi AP to VMs in different availability zones: This last test allows to check whether the connectivity from the Wi-Fi nodes towards any of the service clusters deployed in the city. Again, this experiment allows to compare the nominal capacity of the fibre and copper connections between the Wi-Fi nodes and the destination endpoint. In Barcelona this test was performed, revealing that one of the ports of the Wi-Fi node was underperforming, i.e. delivering a throughput clearly below the 1 Gbps that the port should yield.

Besides testing the virtual networks, the internal throughput for the packets transmitted through the platform should also be assessed. To test this, two VMs should be created as per the tests above:

- two VMs in the edge

- one VM in the edge, another one in the core

Each VM should be connected to both the “management” network and the specific “access” network. The infrastructure provider and the platform maintainer can both connect to the “seed” VM, then manually SSH into the VMs (via the IP given by the “management” network) and then run Iperf commands to determine the bandwidth between the two nodes.

[8] Please note that from i2CAT to the street cabinet a 10 Gbps fibre connection is provided, but the 1 Gpbs port of the edge server is used.

Co-funded by the Horizon 2020
Framework Programme of the European Union

- During this check, the involved actors may find that the throughput is too low. Testing before and after issuing the command "*ethtool -K <interface_name_to_flame_access_network> tx off*" may provide insights on some limitations. Specifically, this command is disabling the checksum offloading; so the procedure is no longer done in the NIC but in the CPU of the VM. This has a considerable negative impact in the bandwidth but, on the other hand, prevents frequent loss of packets that would end up in an unstable environment. This feature must therefore be disabled, at the cost of lower performance.

2.4.6 E2E Readiness Using SFC

The last set of readiness experiments can be described as a "full-stack" approach whereby an SFC is deployed and various KPIs are being measured to conclude the readiness of the platform for experimenters. As part of an internal benchmarking set of tools, FLAME uses a range of Iperf, ping, curl and self-made client/server software packaged into a PeRfOrmance TEst SuiTe (PROTEST). PROTEST comes as a packaged service function with all the software installed and this SF is deployed on each cluster in the platform assigning a unique FQDN. When attaching a client to any AP the full set of KPIs is being tested against all SFEs in the platform.

As the access networks are most likely the bottleneck emulated UEs are being used first as clients which are deployed as a separate stack in OpenStack and attach to the access provider networks of each SR serving an AP. The E2E test using emulated UEs demonstrates the KPIs the platform can achieve which must be close to the numbers obtained with the tests described in Section 2.4.3.

As the next and last step PROTEST is set up on clients that connect over Wi-Fi to repeat the tests using the same locations as for the Wi-Fi tests in Section 2.4.4.

Co-funded by the Horizon 2020
Framework Programme of the European Union

3 PLATFORM DEPLOYMENT

ARDENT, as specified in D3.11, allows the automated deployment of the FLAME platform into an OpenStack-based infrastructure slice. While the FLAME platform in Bristol and Barcelona was deployed via the proof-of-concept implementation of ARDENT, King's College London and Level7 will receive their FLAME platform using the re-engineered solution which offers RESTful service endpoints for the various task of this objective. This section describes the updated workflow as well as the main differences between the prototype and proof-of-concept realisation.

3.1 DEPLOYMENT WORKFLOW

The initial proof-of-concept implementation required to clone the flame-platform repository to a dedicated instance in the infrastructure's environment – either on the controller itself or into a dedicated OpenStack instance set-up for the FLAME tenant specifically. The latter approach in Barcelona already takes away 1 vCPU, 1GB RAM and 50 GB of storage from the slice's resource. Once the repository has been cloned, the proof-of-concept implementation was executed which read in a graphviz[9]-based infrastructure descriptor, performs a sanity check and creates quotas and flavours. Afterwards, the FLAME images were built inside the slice before manually writing the HEAT orchestration template to deploy the platform. With the re-engineered solution, there is no need for any sort of local code clone or image build. Given that the OpenStack controller offers all its APIs as RESTful service endpoints encapsulated in python-based clients, ARDENT can run entirely remotely to create, maintain and delete the FLAME platform in any of the four sites. The process of this activity is illustrated in the figure below.

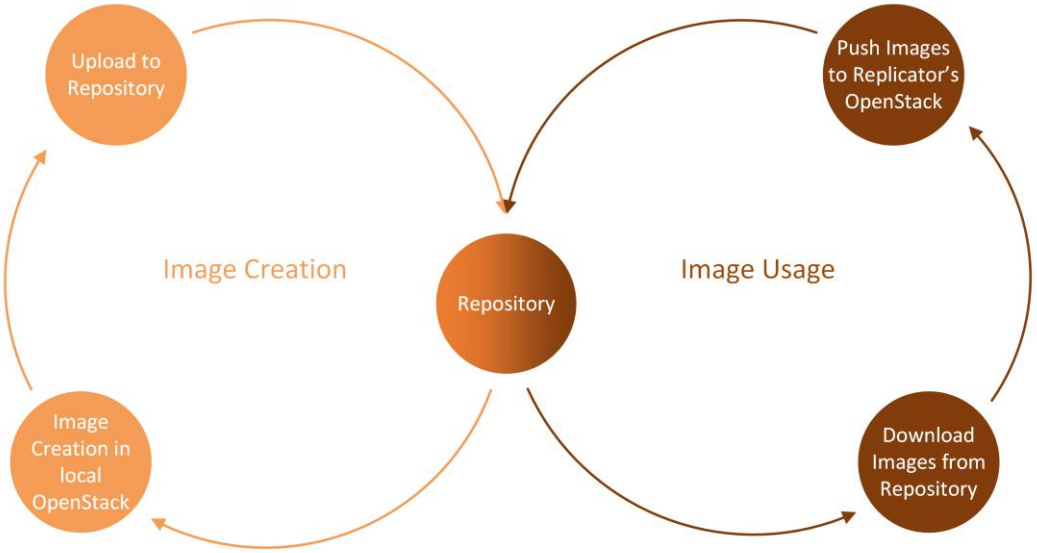

Figure 11: FLAME platform image maintenance.

Figure 11 illustrates a single repository which has the platform images stored. This repository is hosted on IT Innovation's server *givry* which is also hosts to the sandpit. The cloud images for each FLAME

[9] https://www.graphviz.org/

Co-funded by the Horizon 2020 Framework Programme of the European Union

platform component are created in a local OpenStack environment and are uploaded to the repository (left workflow circle in the figure). If a platform at a particular replication site is due to be updated, those images are being imported into the replicator's OpenStack (right workflow circle in the figure).

Once the images have been all imported successfully, ARDENT is being used to upload the infrastructure descriptor which generates the HEAT orchestration template for the FLAME platform at a particular site. The final task is the deployment of the platform by provisioning the HEAT orchestration template (HOT) to the respective OpenStack service endpoint.

3.2 ARDENT SET-UP FOR ALL REPLICATION SITES

In order to support the deployment of the FLAME platform, the software design of ARDENT ensures that the service endpoints provided by ARDENT are decoupled from the state information those endpoints create by means of a database which decouples the service endpoints from any information ARDENT must write. Based upon this chosen software artefact, a single ARDENT node is operating at InterDigital's office with the MySQL database running on the host system and the ARDENT process virtualised into Linux containers, as illustrated in Figure 12.

Figure 12: ARDENT set-up for multi-site deployments.

Each LXC in the figure above has the acronym of the infrastructure provider in its name, i.e.:

- uob: University of Bristol
- i2cat: Fundacio Privada i2CAT, Internet I Innovacio Digital a Catalunya
- kcl: King's College London
- level7: Level7 in Sicily
- soton: University of Southampton (sandpit)
- ide: InterDigital Europe (local staging testbed)

As can be also seen in the figure above, some LXC instances have a VPN box which indicates a VPN client which connects into the replication site allowing ARDENT to reach the OpenStack service endpoints. My moving this up into the LXC potential IP clashes or IP route duplicates across replication sites can be avoided. Each ARDENT process inside an LXC creates a connection to MySQL running on localhost and gets its own MySQL database assigned which keeps the information written by each ARDENT process separate.

Co-funded by the Horizon 2020
Framework Programme of the European Union

3.3 DEPLOYMENT

This section documents the deployment procedures of the FLAME platform using ARDENT. The full API documentation of all ARDENT service endpoints can be found in the appendix of this document. As REST calls are meant for machines to be used, ARDENT comes with a BASH-based CLI which is described first. This CLI builds the foundation for the deployment documentation across the remainder of this section.

3.3.1 ARDENT Command Line Interface

A command line interface is provided which uses *httpie*[10] and *jq*[11] to communicate with all ARDENT service endpoints and parsing the results, respectively. Below, all CLI combinations are listed which are returned when not calling all required arguments correctly:

```
ardent descr add <DESCRIPTOR>
ardent descr udpate <DESCRIPTOR>
ardent descr delete
ardent rc add <RUNCOM>
ardent rc udpate <RUNCOM>
ardent rc delete
ardent hot generate
ardent hot show
ardent hot delete
ardent stack create [--wait]
ardent stack delete [--wait]
ardent stack status
ardent check run [--wait]
ardent check status
ardent check result
```

A detailed description of the arguments above is given below. Note, the reason why some calls allow to provide a --wait argument is that the action to communicate an intent is decoupled from its execution. As given that HTTP by default has a 30 s timeout and it is set to as low as 10 s by the common libraries, the return of the result of the intent could arrive afterwards. Hence, the --wait or a separate status argument is available.

- ardent
    - descr
        - add: Add an infrastructure descriptor to ARDENT
        - update: Update an existing infrastructure descriptor
        - delete: Delete the added/updated infrastructure descriptor from ARDENT
    - rc
        - add: Add the tenant runcom to ARDENT

[10] https://httpie.org/

[11] https://stedolan.github.io/jq/

Co-funded by the Horizon 2020 Framework Programme of the European Union

- update: Update a previously uploaded tenant runcom
- delete: Delete the added/updated tenant runcom

- hot
  - generate: Generate the HEAT orchestration template based on the uploaded infrastructure descriptor
  - show: Show the generate HEAT orchestration template
  - delete: Delete the generate HEAT orchestration template from ARDENT
- stack
  - create: Creating the FLAME platform stack
  - delete: Deleting an existing FLAME platform stack
  - status: Obtaining the status of a stack create or stack delete
- check
  - run: Request to conduct a sanity check
  - status: Obtain the status of the
  - result: Obtain the results of a successfully conducted sanity check

3.3.2 Uploading Infrastructure Descriptor and OpenStack Runcom

First, the OpenStack tenant runcom is uploaded to ARDENT as well as the infrastructure descriptor. Both files are provided by the infrastructure provider.

```
root@ardent:~/# ardent tenantrc add ~/kcl-tenantrc
---
Adding new tenant runcom
---
{
  "status_id": 0,
  "status_str": "Request Successful"
}
root@ardent:~/# ardent descr add ~/kcl-descr.yaml
---
Adding new descriptor
---
HTTP/1.1 200 OK
Content-Length: 50
Content-Type: text/plain; charset=utf-8
Date: Sun, 20 Oct 2019 14:50:11 GMT
{
  "status_id": 0,
  "status_str": "Request Successful"
}
```


Co-funded by the Horizon 2020 Framework Programme of the European Union

3.3.3 Sanity Check

Next up, ARDENT must be instructed to perform a sanity check to ensure the necessary OpenStack configurations have been completed (see D3.11 for details):

```
root@ardent:~/# ardent check run –wait
---
Requesting sanity check
---
HTTP/1.1 202 Accepted
Content-Length: 0
Date: Sun, 20 Oct 2019 14:44:44 GMT

---
Waiting to complete.....................................
---
HTTP/1.1 200 OK
Content-Length: 54
Content-Type: text/plain; charset=utf-8
Date: Sun, 20 Oct 2019 14:45:49 GMT

{
    "status_id": 0,
    "status_str": "Sanity-Check completed"
}
```

3.3.4 Generate HEAT Orchestration Template

After ARDENT has performed its sanity check, the HOT can be generated following the assumptions about the compute nodes and which FLAME they will host outlined in D3.11.

```
root@ardent:~/ardent/cli# ./ardent.sh hot generate
---
Generating HOT
---
HTTP/1.1 200 OK
Content-Length: 50
Content-Type: text/plain; charset=utf-8
Date: Sun, 20 Oct 2019 15:02:14 GMT

{
    "status_id": 0,
    "status_str": "Request Successful"
}
```

3.3.5 Stack Deployment

After successfully completing the steps above, the FLAME platform stack can be deployed:

```
~$ ardent stack deploy --wait
---
Deploying stack ‘flame-platform’
---
HTTP/1.1 202 Accepted
Content-Length: 0
Date: Sun, 20 Oct 2019 15:10:53 GMT
```


Co-funded by the Horizon 2020 Framework Programme of the European Union

```
Waiting to complete.....................................
HTTP/1.1 200 OK
Content-Length: 54
Content-Type: text/plain; charset=utf-8
Date: Sun, 20 Oct 2019 15:15:49 GMT

{
    "status_id": 0,
    "status_str": "CREATE_COMPLETE"
}
```


Co-funded by the Horizon 2020
Framework Programme of the European Union

4 SPECIFIC DEPLOYMENT INSIGHTS

In spite of a meticulous planning of a FLAME deployment as it has been done for the replicator sites, there is always a chance that some conceptual details in practice turn out to be different. This can lead to unexpected issues that are determined either in the validation phase or even later, during the operation of the FLAME infrastructure. While the rest of this document already points out what is required to successfully replicate FLAME, this section gathers information that is based on specific experiences in the replicator sites and may prove useful for any future replicators.

4.1 BRISTOL DEPLOYMENT INSIGHTS

Access to the Millennium Square

Experimentation in the Bristol Millennium Square needs permission from the City Council. Not only the trial should not clash with other events but also the Council should be aware of the experimentation footprint and the way it affects the public using the space. A data sharing agreement template consistent with GDPR is available, agreed between UoB and each experimenter.

Network Sanity

Using a tool developed by Interdigital called PROTEST, it has been possible to locate several data bottlenecks. In one case the packets dropped significantly between the Edgecore and the datacentre. The case was escalated to the Pica8 support that after analysing the logs identified wrong settings in the Edgecore switch. The use of this tool is therefore highly recommended to assess the network health, well in advance to an experimentation campaign.

Wi-Fi congestion

Besides the FLAME platform, the Wi-Fi radio interface has limited resources. During a major event when 20 users were uploading huge data on FLAME platform on Wi-Fi using laptops, phones etc., the Ruckus controller raised flags of congestion. There was a bottleneck of the number of allowed users that was causing the alarm.

Concurrent use of trial space

While Millennium Square has many benefits, it can sometimes be blocked to experimenters due to clash of events. Therefore, as a redundancy plan, a pedestrian path just 300 meters further down, has been added to the FLAME network. This area is covered by two APs which are also fibred to the compute servers at the Square. This area will be used as plan B, when testing in Millennium Square is not possible.

Broken SDN fabric

At one occasion one of the compute nodes lost power and a fibre link was broken. This came into attention only when FLAME was being tested. The lack of any FLAME SSID raised the alarm and triggered prompt action. However this could have been addressed much earlier if the SDN topology was checked regularly (see Section 2.3.1.3), or a monitoring tool could have identified the broken link real-time.

Co-funded by the Horizon 2020 Framework Programme of the European Union

4.2 BARCELONA DEPLOYMENT INSIGHTS

Cooling of Wi-Fi Nodes

For Barcelona a custom Wi-Fi solution was designed and built. Once the prototype of the i2CAT Wi-Fi nodes was finalized, we detected that the temperature inside the casing would grow quickly during the operation of the Wi-Fi cards. Part of this issue was caused by the great amount of heat the Wi-Fi interfaces would generate (up to 3 transceivers per Wi-Fi node) during operation. Tests in the lab revealed that temperatures in a Wi-Fi "box" would rise quickly once one or several interfaces were transmitting data and this was measured while in a room with around 23ºC. Considering the much higher environmental temperature we could get on-street, a redesign was necessary. Experiments in the lab revealed that installing a small fan on both sides of a Wi-Fi box would generate an airflow sufficiently strong to cool down the boxes' interior. As such, all Wi-Fi boxes were redesigned with the approval of the manufacturer (to assure that the IP rating of the boxes are maintained) and the improved version was eventually mounted on the lamp posts. During the whole operation of FLAME in Barcelona no further temperature issues were detected in the Wi-Fi nodes.

Cooling of the cabinet (edge compute)

After deploying the edge compute in the cabinet, we noticed that the server in the cabinet could reach high temperatures during summer, especially the regular heat waves that are prone to happen in Barcelona. Anticipating to that situation we decided to turn the servers off prior to episodes of high heat. To this matter, we recommend doing some initial stress tests and ideally some remote test to determine whether the IPMI is enabled and the correct status of the link connecting the server and the local site from where to issue the remote shutdown/boot commands take place. Further, to lower the temperature in the cabinet in more general terms, a solution was installed that would suck cool air from the underground and blow it into the cabinet. This reduced the peak temperatures by around 10ºC, making it possible to run the equipment even in summer (except for aforementioned heat waves).

Public space for experimentation

The location of the nodes, ideally, should have nearby public benches and/or restaurants, coffees or other shops where it is possible to conduct experiments on-site with a strong Wi-Fi signal. This will have a positive impact on the bandwidth and provide for a better Quality of Experience (QoE) on the media transmission. In the Barcelona deployment the conditions in general are favourable in this sense, as there are several bars with outdoor tables and chairs and there are many benches and sidewalks/traffic reduced streets to conduct experiments. Still, in colder periods of the year when sitting outside is not possible or whenever it rains, the indoor coverage obtained in nearby bars and restaurants is not good enough to conduct experiments.

Using Wi-Fi in the unlicensed band

Since the Wi-Fi APs operate in the unlicensed Wi-Fi band, there is always a risk of experiencing connectivity issues due to the interference. While it is commonly possible to avoid crowded channels, the choice of Wi-Fi channels for FLAME is limited in several ways:

- the use of the lower 5 GHz band is required, since we only can use the 5 GHz channels and most of the UEs do not support the upper 5 GHz band

Co-funded by the Horizon 2020 Framework Programme of the European Union

- Choosing 40 MHz bandwidth for the channels results in only two channels being available in the lower 5 GHz band: channel 36 and channel 44. While there are more options using 20 MHz channels, the delivered throughput is not good enough to serve multiple users at once on a lamp post. Also, while 80 MHz channels can deliver very high throughput, there is only one such channel available in the lower 5 GHz band and it is likely for this channel to be interfered by users operating in any of the other, non-FLAME Wi-Fi APs.

As such, while we can reduce the cross-lamp post interference by assigning different channels to neighbouring lamp posts, the FLAME Wi-Fi is exposed to interference of any other Wi-Fi in the area.

Delays in getting permissions and procurement

Execution of constructions and getting permissions can take longer than anticipated, especially when deploying in public spaces and when needing approval from city council. This can lead to delays in the timeline of a FLAME deployment, as it has been experienced in Barcelona. Further, for some construction orders or to contract services (public procurement) certain time margins have to be assumed.

Neighbour complaints (radio deployment)

It can happen that installing radio equipment with visible antennas (omnidirectional/directional) as it has been done in Barcelona can cause neighbours to raise complaints, as they feel that the radiation of the equipment could be noxious or even the integration of the elements in the urban furniture has a negative impact on the visuals of the urban environment. In Barcelona, the risks for such complaints were minimized by covering most of the equipment inside a *Radomo*[12] that nicely blends in with the street furniture. Further, the Wi-Fi equipment was configured to operate with low power (unless higher powers are requested by experimenters) to reduce the real impact Wi-Fi radiation might have on the environment and as a preventive action in case neighbours complain. In spite of these preparations, complaints were raised during the operation of FLAME and the operation had to be stalled for some time, resulting in a downtime for the FLAME infrastructure. After presenting the facts to the neighbours (low transmission power, equipment is only used sporadically), the operation was reassumed.

4.3 BUSETO PALIZZOLO INSIGHTS

Cooling of the cabinet (edge compute)

For some nodes, the installation of the edge compute is done in the outdoors, which can cause issues during the summer time. For this reason, Level7 is installing new cabinets with expanded room and fans in order to keep the temperature of the devices as low as possible. It also must be noted that the SDN OpenVswitch has been implemented with embedded hardware that has an extended temperature.

[12] A cylindrical construct enveloping the Wi-Fi nodes that integrates with the street furniture (lamp post)

Co-funded by the Horizon 2020 Framework Programme of the European Union

Experimentation in public areas

The infrastructure is owned and operated by Level7 but it is also installed in public places, in order for the community to take advantage of the services and support experiments during public events or in public areas.

For example, the infrastructure has been deployed in the local museum, that hosts objects from the cultural heritage of the rural community. Therefore, the access to some places (library, school, museum, etc.) must be agreed in advance with the local community in order to carry experiments in those places.

For other nodes, the coverage is provided in public places that are fully open to the public, such as the square or other outdoor places. In this case, any activity such as experiments can be done without any special “assistance” or permission to be asked in advance.

Using Wi-Fi in the unlicensed band

Since the access network is operated in the unlicensed Wi-Fi band (2 GHz and 5 GHz), there is always the risk of experiencing frequency overlapping/interferences due to other devices in proximity of FLAME device.

It also must be noted that this can happen more frequently only for some nodes of the infrastructure, while other nodes are deployed in the indoor where the walls will provide a natural shield for interferences. For those indoor scenarios, 40 MHz and 80 MHz channels can be provided in order to fully take advantage of the provided bandwidth.

4.4 LONDON INSIGHTS

As we had the infrastructure readily available and used by other third-party vendors, insights from physical deployment perspective is limited. For example, the installations and regulations were already in place. Thus, our major insights come mainly from the difficulties of having a collocated edge services, where there is a requirement for FLAME to exist with other vendors sharing the infrastructure. In the following we describe the 2 major challenges that had to be overcome:

Security Plugin and Security Groups

We had to enable port security in our OpenStack, which we did not support to our other tenants. We noticed that the ARDENT tool checks whether the plugin is loaded rather than checking the individual ports. Again, we had to enable security group just for FLAME, however the firewall component is disabled on all hosts so any configuration will be gracefully ignored.

Admin Rights

Admin rights in OpenStack are mainly required by FLAME for following reasons:

- stack deletion
- create/modify and delete flavors within OpenStack
- hostname is otherwise not visible as a non-admin

Co-funded by the Horizon 2020
Framework Programme of the European Union

We also noticed that some flavours of other tenants on the KCL testbed got deleted accidentally, thankfully not resulting catastrophic failures as we had a backup of the flavours. Admin rights also gives the visibility of other tenants which is not suitable for large scale deployments.

Various ideas were brought out to improve FLAME and allow the platform to co-exist within an operator's infrastructure. A simple solution we adopted is to physically isolate FLAME from the more sensitive components, putting them on a different machine. Other solutions include VLAN-level isolation etc. Thus, by deploying FLAME on a testbed with other applications, the KCL replication helped refine and make explicit the requirements of the infrastructure (in terms of access to very components within the infrastructure) and to standardise the edge requirements for orchestration.

Co-funded by the Horizon 2020
Framework Programme of the European Union

5 ASSURING EXPERIMENTERS' READINESS

5.1 EXPERIMENTER'S WORKFLOW: FROM THE DESKTOP TO USER TRIALS

FLAME is enabling trials of innovative media applications with people in real-world locations. By creating and supporting a process to take ideas from the developer's machine to user trials in a FLAME city or replica, we have reduced the cost of testing new services and experiences with the public. This process is FLAME's devops pipeline.

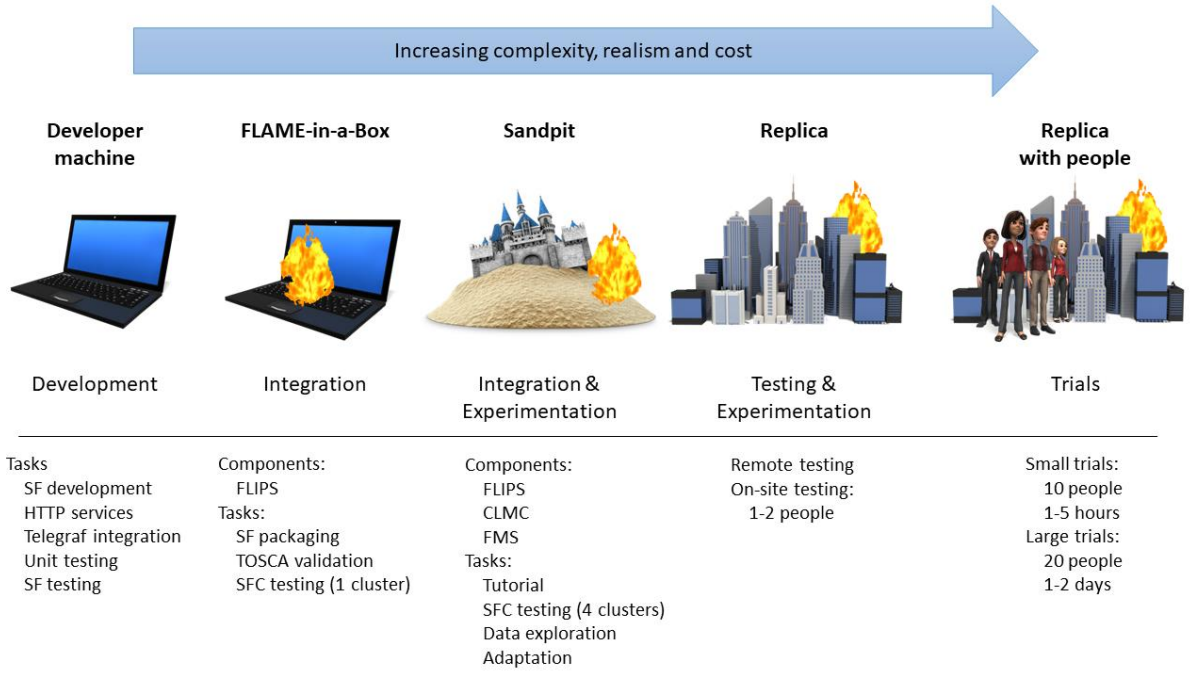

Figure 13: FLAME's devops pipeline showing the activities at each stage.

The pipeline (Figure 13) takes software from the development machine to deployment for user trials in a FLAME replica such as Bristol or Barcelona. Going from one end to the other increases the complexity, realism and cost, moving from TRL 4 (technology development) to 6 (technology demonstration). Each stage must be completed before moving on to the next, but it is expected that earlier stages will be revisited as lessons are learnt and ideas and concepts are revised and incorporated (Figure 14).

Co-funded by the Horizon 2020 Framework Programme of the European Union

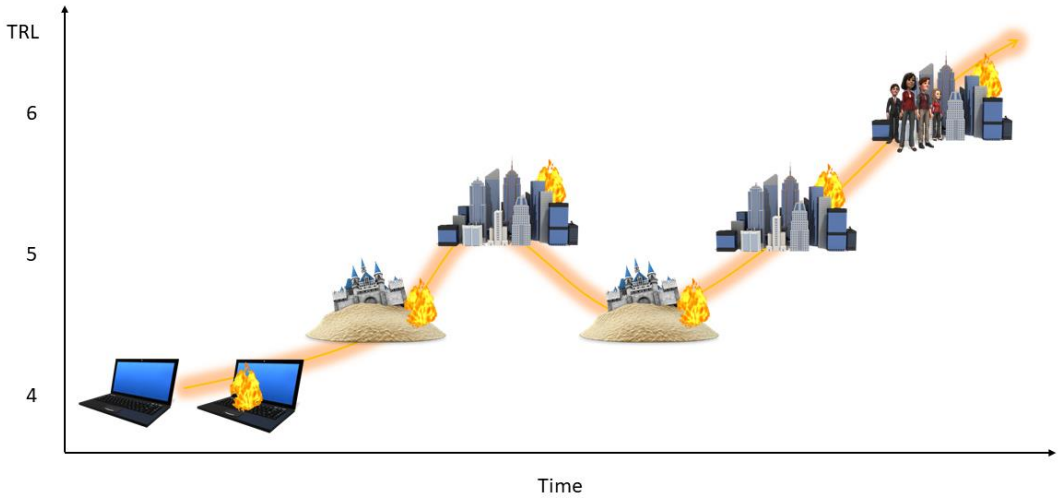

Figure 14: A development process will repeat earlier stages as lessons are incorporated. The TRL increases from the start of the pipeline to the end.

5.2 DECOMPOSITION OF THE SERVICE AND INITIAL DEVELOPMENT

The initial development of a (media) service often involves the decomposition of an existing software service to create several independently operating service functions (SF) that are communicating over the network with each other ideally over HTTP and composed into a service function chain (SFC) to create the whole service. Once the decomposed service's unit tests have passed it is packaged for FLAME using a packing toolchain provided by the project for the LXD and KVM hypervisors. As part of the packaging toolchain, Telegraf is installed which allows the collection of various performance data points of the operating system the service container and the service itself. The result of the packaging is a compressed TAR ball which can be uploaded into any FLAME-based platform.

5.3 INITIAL INTEGRATION USING FLAME-IN-A-BOX

FLAME provides a virtual appliance, FLAME-in-a-Box, that fits on an ordinary modern office laptop. As illustrated in Figure 15, this appliance comprises a minimal set-up of virtual instances allowing experimenters to familiarise themselves with the TOSCA-based resource descriptor and test it:

- a UE (user equipment node) for the test or client software;
- a "cluster" where packaged service functions are deployed;
- the "sr-ue" which is a SR connecting the UE, cluster and pce-sfemc;
- the "pce-sfemc" node (path computation element and service function endpoint management and control services) which also includes the FLAME orchestrator;
- another SR ("sr-ps"); and
- a "ps" instance for platform services such as DHCP, IP gateway and DNS.

Additionally, the FLAME orchestrator node has been also equipped with a web-based TOSCA parser allowing experimenters to get some help in case they have written an invalid TOSCA descriptor.

Co-funded by the Horizon 2020
Framework Programme of the European Union

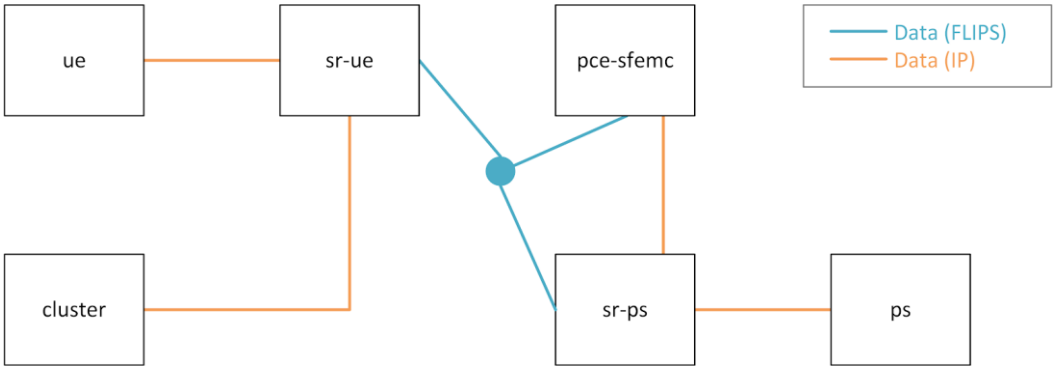

Figure 15: FLAME-in-a-Box setup on a local machine.

5.4 INTEGRATION & EXPERIMENTATION IN THE SANDPIT

FLAME's sandpit is a 72-core server hosted at IT Innovation and accessed remotely through SSH and web interfaces. The sandpit uses a combination of containers and virtual machines to emulate a deployment of the FLAME platform in a physical infrastructure such as that found in Bristol's Millennium Square.

In contrast to FLAME-in-a-Box, the sandpit provides FLAME's cross-layer management and control (CLMC) component and sufficient resources (memory, CPU) in four "clusters" for the deployed service functions to execute and be tested. Crucially, the sandpit also includes four "emulated UE" nodes which allow experimenters to install their test clients on the user-equipment (UE) nodes which can then connect to and test their service functions. Each UE is deployed in a different area of the sandpit to be closer to a particular cluster and thus vary their performance (Figure 16).

The sandpit also hosts FLAME's interactive tutorial which guides a developer through creating a TOSCA resource specification and TOSCA alerts specification, deploying a simple service function chain, viewing its state in the orchestrator, viewing monitoring data in the CLMC, and using one of the UEs to apply load to the service and triggering a change in the deployment.

Co-funded by the Horizon 2020
Framework Programme of the European Union

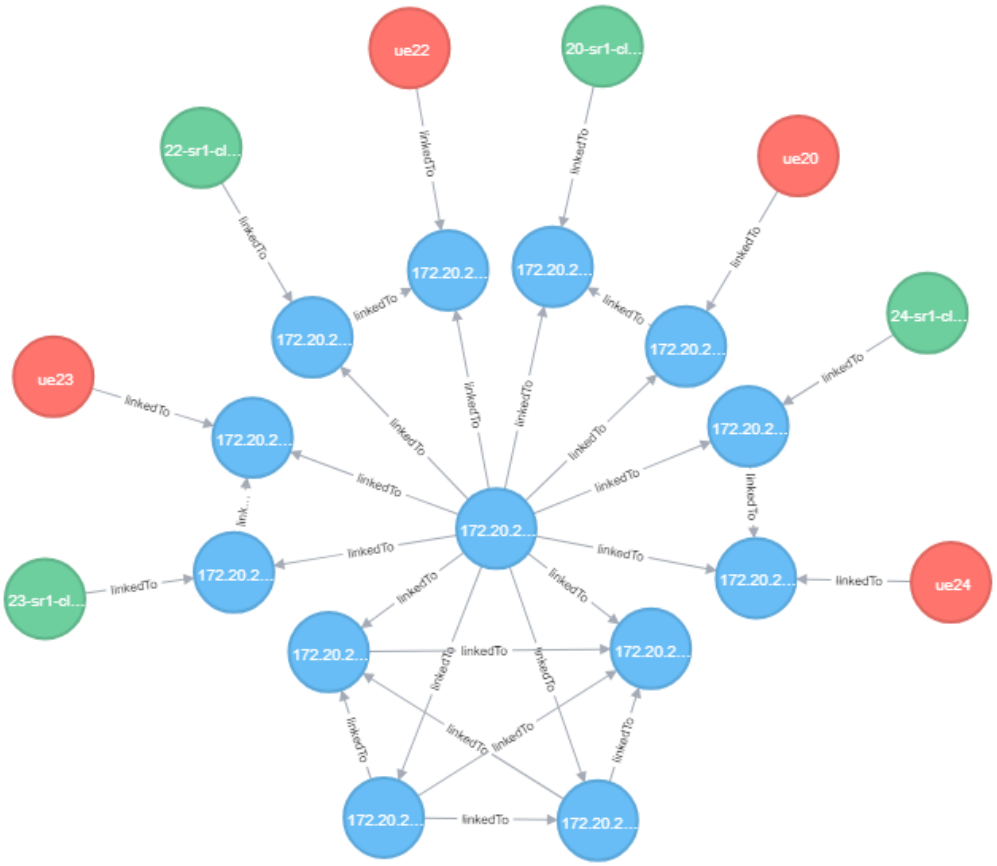

Figure 16: The topology of the sandpit with clusters (green), emulated UE (red) and SDN switches (blue).

In the sandpit, the media service developer uploads their packaged service functions to the sandpit's repository and adapts their service function chain (TOSCA resource specification) to make use of the four available clusters and, potentially, the foundation media services which are provided by FLAME and made available through the same image repository.

The key tasks in the sandpit are to test the functionality of the deployed service function chain (but not test its performance) and understand and explore the data collected in the CLMC. Data is collected at various levels: metrics from the SDN level describing routes and latencies; metrics from the container of each deployed service function (service function endpoint or SFE) describing the memory, disc, CPU and network; metrics describing the performance of application containers or web servers such as Tomcat or Nginx if they are used; and application-specific metrics where they are implemented. This data may be explored in the CLMC's Chronograf interface and used to understand the behaviour of the service and to define triggers (encoded in the TOSCA alerts specification) which fire off alerts to the SFEMC or to the media service in order to adapt the service to demand. Note, while the TOSCA-based resource descriptor is given to the FLAME orchestrator, trigger alerts are sent to the SFEMC, as both are operating at different timescales.

Functional tests and experiments to understand a service's response to demand may be conducted from the emulated UEs. A developer uses SSH to get a terminal on a UE and may install their client software or test suite there. It is even possible to forward an X-server connection from a UE to a developer's desktop if a graphical interface is required.

Co-funded by the Horizon 2020
Framework Programme of the European Union

5.5 TESTING AND EXPERIMENTATION AT A REPLICA

Although the sandpit provides a realistic technical environment, moving from the sandpit to a replica does introduce key changes. The performance (network, CPU) in a replica will be different to the sandpit which affects the performance and behaviour of the media service. The physical environment provides new challenges: the connectivity of mobile phones must be tested (both the signal strength, bandwidth and which AO is used in different locations) and the behaviour and positioning of physical markers (e.g. QR codes, Bluetooth Low Energy beacons) must be experimented with.

Ideally, the first tests are done remotely so that the developer does not need to travel. A VPN connection to the replica can be used to deploy a service function chain and some tests can then be performed either by including an emulated UE as part of the chain or by a member of the FLAME team local to the deployment physically performing some tests using a mobile phone under the instruction of the developer. Data can be explored in the remote replica's CLMC interface.

As for the order in which emulated and physical UEs are used for experiments, the experimenters are encouraged to use emulated UEs first and as such to test their SFCs remotely without the need of on-street support. This step allows the experimenter to get a general idea of the performance in the infrastructure (obviously without taking into account the radio access network) and it allows to test at least a subset of an SFCs function (as some tests may require physical UEs to be used). Whenever tests with emulated UEs are concluded, i.e. all desired and possible tests have been performed, or when an experimenter would require the use of real UEs, the experimenters would move on to doing experiments with on-street support.

Preparing on-street experiments with FLAME team members under an experimenter's remote supervision has a higher organizational overhead, when compared to the previous steps in the experiment workflow (e.g. reserving a time slot in the Sandpit and being assigned the required resources or reserving the infrastructure for remote access and emulated UE testing):

- Assuring that FLAME team members of the infrastructure provider are available (number of users to be specified by the experiment)
- Assuring that technical support from team members responsible for FLAME platform components is available
- Reserving a slot in the infrastructure and assuring the infrastructure's operation and maintenance teams are aware of the tests to be performed
- Gathering the required UEs to do testing and with the required settings (Smartphones, tablets, laptops with specific tools, browser versions, operating systems, etc.)
- Handing over an accurate plan of the experiments to be done to assure maximum efficiency of the time spent at the deployment

Generally speaking, each experimenter should be assigned at least 2 days on which they can perform tests with real UEs and on-street support. As experiments progress, experimenters and on-street support team evaluate when the experiment is working fine and a day for a demo trial is chosen. On the day of a demo trial, the on-street support team mimics all the details of the final trial, e.g. number of users, numbers and types of UEs, workflow of the experiments, events happening on-street during an experiments and so on. This step helps to validate that the experiment design, but also the trial plan can be executed without any issues. Of course there are limits to the number of users that can be available on such a demo trial, but the intention is to get as close as possible to the real number of users that would be on street during a real trial.

Co-funded by the Horizon 2020
Framework Programme of the European Union

With these remote tests successfully completed, a developer (with the agreement of the host) would travel to the replica site to perform those tests and experiments which are dependent on the physical environment. This is also an opportunity to get an initial impression of the user experience in the environment.

5.6 USER TRIALS

The ultimate aim is to understand the user experience of a new media service and how FLAME's innovative features serve to enhance that experience. This is achieved through the user trials. Before conducting user trials there are many administrative tasks to be completed such as participant recruitment and the preparation of data sharing agreements, participant information and consent sheets and so forth but, whilst important, they are not discussed further here.

Small Trial

A "small trial" may be around 10 people on site over some short time and might not include all aspects of a media service (it could just focus on part of the experience).

Small scale FLAME trials serve to provide a number of important outcomes:

- validation of the technical deployment and operation of your media service;
- generate observations of real use that will drive usability improvements; and
- capture of a selection of QoE data that can be analysed alongside QoS metrics.

If well-planned, small scale trials will generate critical data (and later, knowledge) that will significantly improve the quality of a FLAME based media service as well as removing or mitigating risks before scaling to a large trial.

A small trial should include time for the participants to reflect on the experience. At its most basic, participants should be asked to respond to a simple questionnaire. However, greater insights can be garnered by actively reviewing experiences with your participants – for example, by discussing particular observations you have made and asking further questions to clarify behaviours.

Feedback from a small trial may cause a developer to rethink the service and its experience and so it may be necessary to return to the sandpit to understand how to capture and respond to additional data or back to the first stage to adapt the service and repackage.

Large Trial

A "large trial" would be around 50 people on site, free to use the service as they like. Primary outcomes for a FLAME large trial are:

- verified, stable deployment and provision of your media service at scale;
- demonstrable, dynamic management of your media service using FLAME platform behaviours; and
- empirically supported knowledge supporting the benefits of the FLAME platform and the value of the media service.

The scale of a large trial causes differences in the approach and challenge in comparison to a small trial. For instance, with many people it is not possible to use qualitative analysis methods (such as individual interviews) to gather information about the users' experiences. The scale of the gathered

Co-funded by the Horizon 2020 Framework Programme of the European Union

data set is dramatically increased and, combined with the less-controlled participant behaviour, this may become difficult to analyse.

Opportunities come from having a sufficiently large cohort to do effective A/B testing (if appropriate) and to gather a dataset which can provide significant statistics and evidence for the efficacy of the approach.

Co-funded by the Horizon 2020
Framework Programme of the European Union

6 REPLICATOR PROGRESS

Although the Bristol and Barcelona replicas have been supporting experimentation for a long time now, they have been further developed to improve performance and offer more possibilities to experimenters. Processes have been refined, the FLAME platform software updated frequently and topologies changed and extended. Indeed, the development of the platform and the replicators in response to need is an integral part of providing and experimental and experimentation platform.

The latest two replicators, London and Buseto Palizzolo, have been working closely with FLAME, following the replication process described in the first version of this document (D5.1) but with the updated technical knowledge and procedures described above. The replicators started work in May 2019 knowing they must be prepared to accept experimenters (from open call 3) from January 2020 onwards. Status updates on the technical readiness may be found below.

In addition to the technical aspects, both new replicators have made sure to address the business aspects of the replication process. They have looked at and described in their reporting the stakeholders, financing, regulation and certification, roles and governance.

6.1 OVERVIEW OF PLATFORM REPLICATION SITES

FLAME offers four replicas to experimenters, as described in the introduction. In addition to that the DevOps workflow designed for experimenters includes FLAME-in-a-Box as well as the sandpit which is made available to FLAME validation and open call experimenters. With the help of ARDENT to fully automate the platform deployment in OpenStack-based environments it becomes a scalable model to deploy different (new) platform versions across multiple sites almost at the press of a finger..

However, FLAME can also be deployed using a toolchain developed by InterDigital Europe to bring the service function routing software already deployed Linux machines (either virtual instances or native deployments) in an automated fashion which installs the software and configures it according to the environment. The combination of ARDENT and the deployment toolchain for native Linux environments allows the automation of the FLAME platform deployment at the Open Call 2 winner Level7 in Sicily.

Given the flexibility of deploying the platform the map in Figure 17 illustrates the locations where FLAME is deployed. Besides the known four replicator locations, the sandpit and FiaB across tens of laptops, InterDigital has it deployed inside their staging testbed (OpenStack based), their profiling testbed LEARNT as well as in a mobile demo aimed at exhibition events such as Mobile World Congress. In addition to that, the SFR stack is deployed in Conshohocken in InterDigital's former lead developer's lab environment for development support coordinated out of the London office. Lastly, the platform is also deployed at InterDigital's Berlin office the 3GPP SBA efforts.

Co-funded by the Horizon 2020
Framework Programme of the European Union

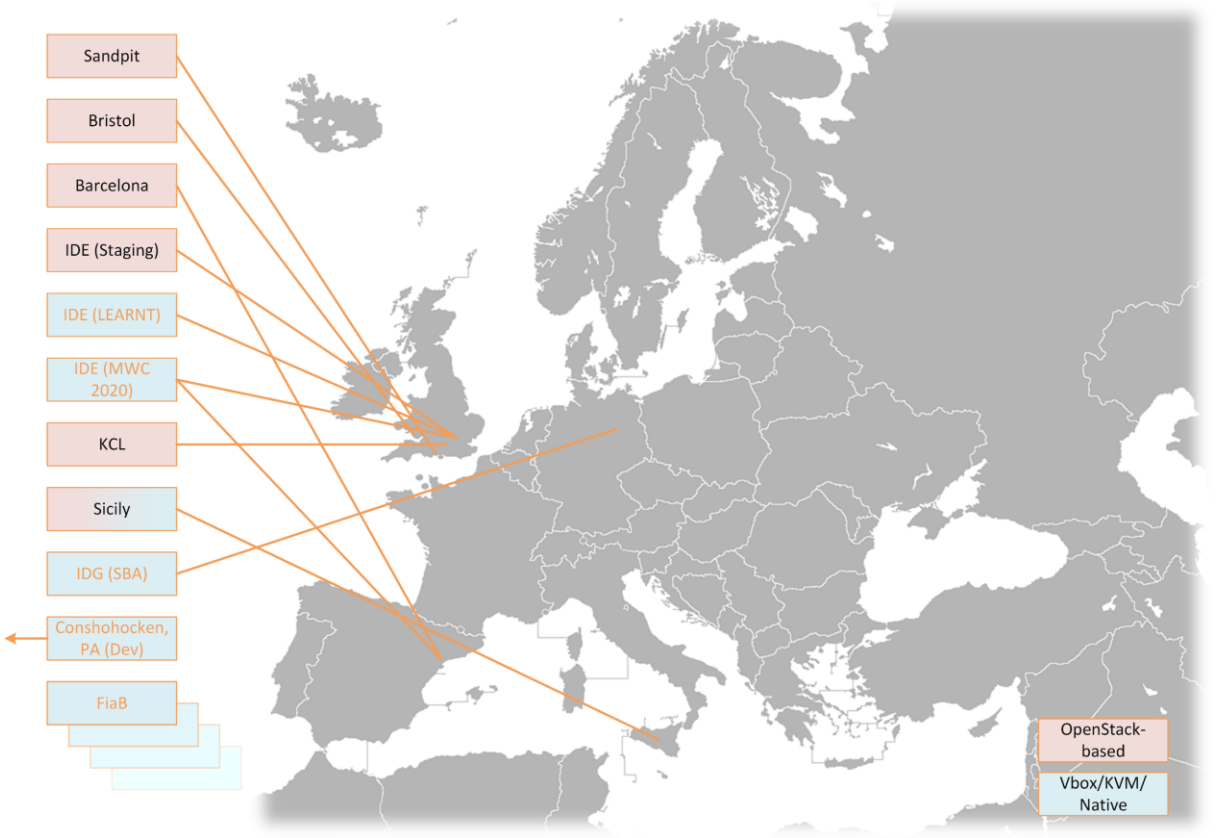

Figure 17: Map of deployed platforms.

6.2 BRISTOL

Over the course of the last 12 months there were several occasions when the Millennium Square was reserved for non-FLAME events and planned trials could not take place. As a contingency plan UoB was able to launch another FLAME "open" testbed just 300 meters from the Square which has been fully integration into FLAME. This new site, opposite the M-Shed museum, offers two Wi-Fi access points which are covering different areas of the museum in comparison to the Millennium Square.

Another major improvement of Bristol's infrastructure has been achieved with the insights from OpenCall 1 trials and the hackathon where under certain load scenarios a significant performance degradation had been observed. Further investigations revealed packet drops by the SDN hardware switches. Thanks to maintenance contract with Pica8, the vendor of the operating system PicOS, the limitations of the TCAM size in respect to the interpretation of the FLAME forwarding rules was revealed and addressed by updates to the operating system and changes to the rules generated by the FLAME platform.

6.2.1 Infrastructure Slice

With the addition of the location M-Shed the infrastructure slice has been extended, as illustrated in Figure 18. The two new Wi-Fi access points can be seen in the bottom right corner, labelled as FLAME-ME (for M-Shed East) and FLAME-MW (M-Shed West). As there is no additional compute node added to the new location, the traffic arriving from those access points are switched up to Tower 2's compute which had only been used for steering traffic from the single SSID that covers the entire square and

Co-funded by the Horizon 2020 Framework Programme of the European Union

provides L2 mobility transparently to the UEs and the platform. Additionally, the RAM of the dc8 compute has been doubled from 32GB to 128GB.

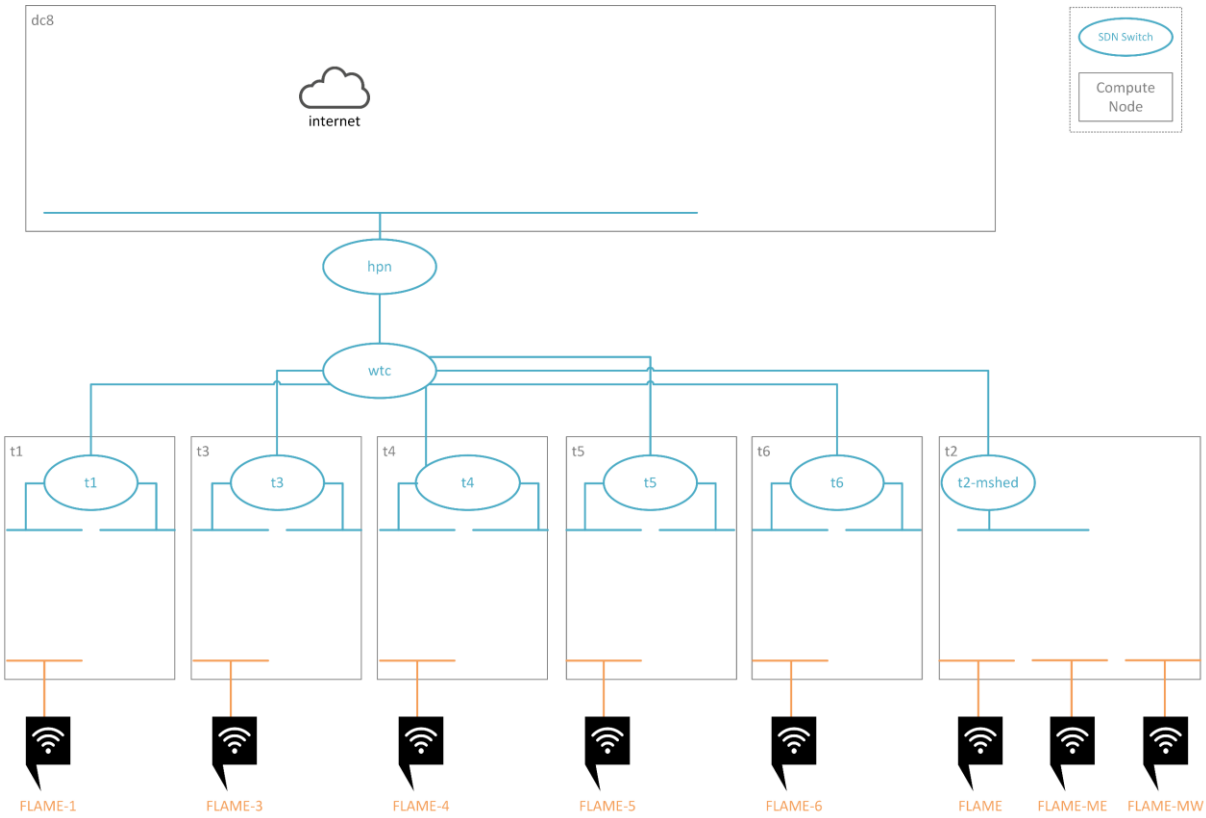

Figure 18: Infrastructure slice for FLAME in Bristol.

6.2.2 Platform

With the addition of the two access points at M-Shed to the infrastructure two new SRs have been added to Tower 2's compute node, as illustrated in Figure 19 at the bottom right corner. Therefore, the cluster instance at compute node t2 is now serving users attached to the single SSID as well as M-Shed users in case the appropriate SFEs had been deployed there as part of an SFC. With the fixes to the SDN hardware switches described in the introduction of this section, the wtc EdgeCore switch now allows to offer a dedicated 1G port to each tower with all SDN rules being executed using the TCAM.

Additionally, with the increase of RAM in the dc8 compute node where all dc8-* instances are deployed experimenters have 74 GB of RAM available I their dc8-cluster for SFEs.

Co-funded by the Horizon 2020 Framework Programme of the European Union

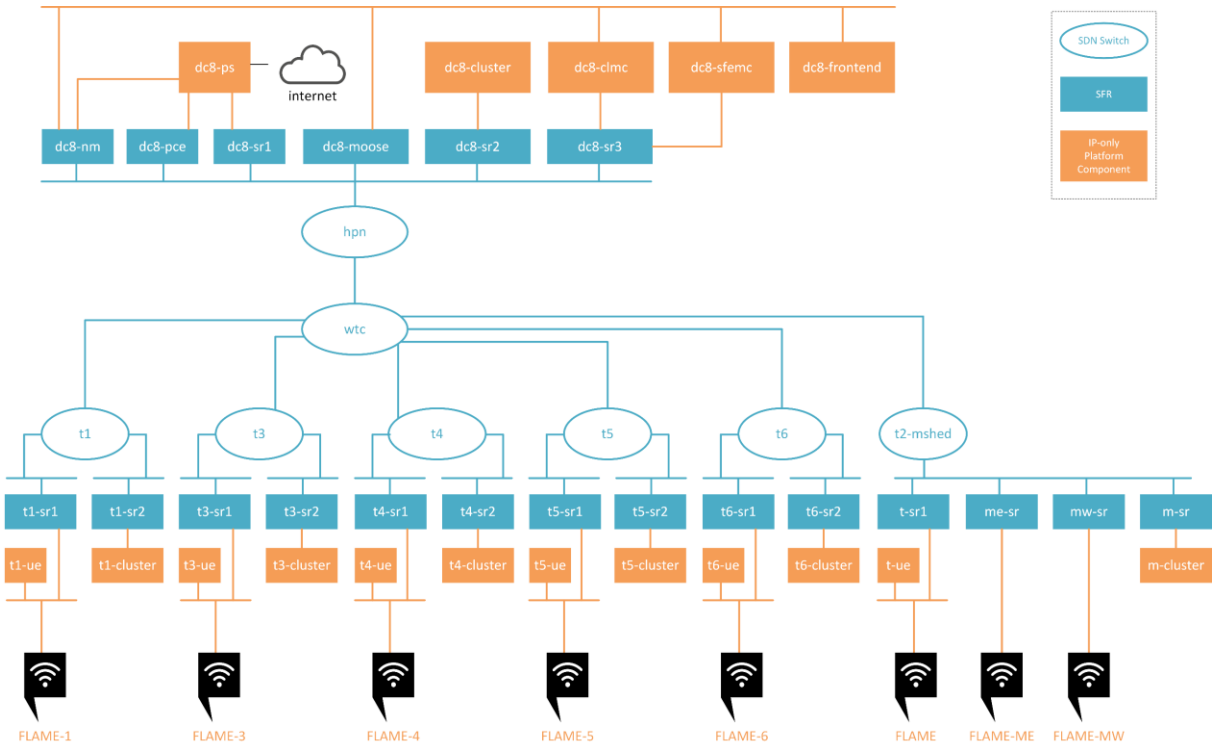

Figure 19: Platform topology in Bristol.

6.3 BARCELONA

For open call one the platform deployed in Barcelona offered a single large cluster at the edge allowing experimenters to bring their SFs closer to the user. Also, all SRs serving the Wi-Fi access points were interconnected via a Wi-Fi-based multi-hop backhaul which did not provide the performance expected. Thus, as a first step the wireless backhaul had been removed, as it did not add anything to the deployment in regards to the resulting topology and its ability to foster a better edge deployment. In addition to that, it was understood from the first set of experimenters that the availability of a single cluster at the edge serving did not allow the same flexibility for the orchestration of SFEs as Bristol offered where a cluster at each access point is present. This has been addressed by following the same methodology even though the size of each cluster is rather limited in terms of cloud resources (CPU, RAM, disk).

6.3.1 Infrastructure Slice

The infrastructure slice in Barcelona for the FLAME project has not been changed significantly in comparison to when the wireless backhaul was in place. Very similar to Bristol, a dedicated SDN switch is present at each Wi-Fi access point location. What is important to highlight though is that a single compute node, fog1, offers access to all four SDN software switches and the access networks.

Co-funded by the Horizon 2020 Framework Programme of the European Union

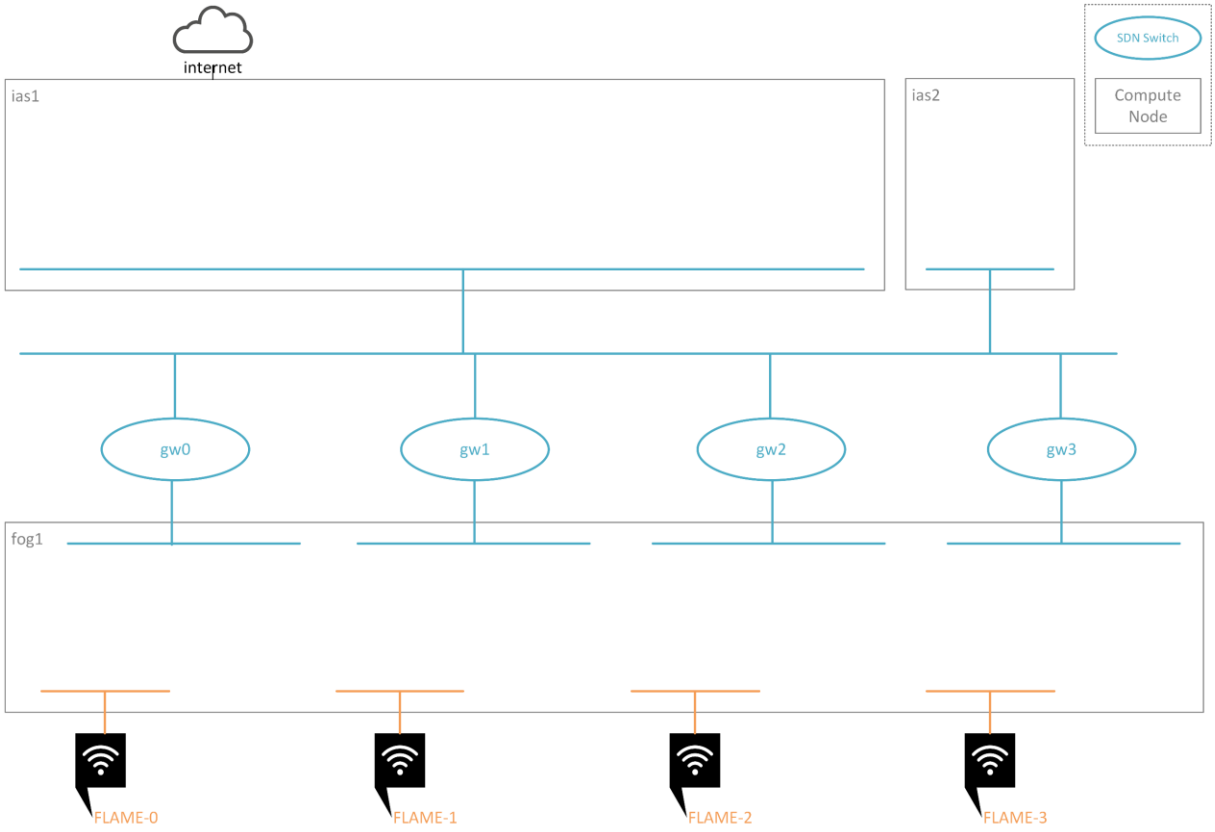

Figure 20: Infrastructure slice for FLAME in Barcelona.

6.3.2 Platform

With the availability of all four access (Wi-Fi) and data (SDN) networks at the single edge compute node, the platform offers now a dedicated cluster at each access point, as illustrated in the figure below. Even though each cluster providers rather low compute resources it allows experimenters to have SFEs right at each access points which is more important to utilise the advantages of the service function routing capabilities of the platform than a larger cluster serving all four locations.

Co-funded by the Horizon 2020 Framework Programme of the European Union

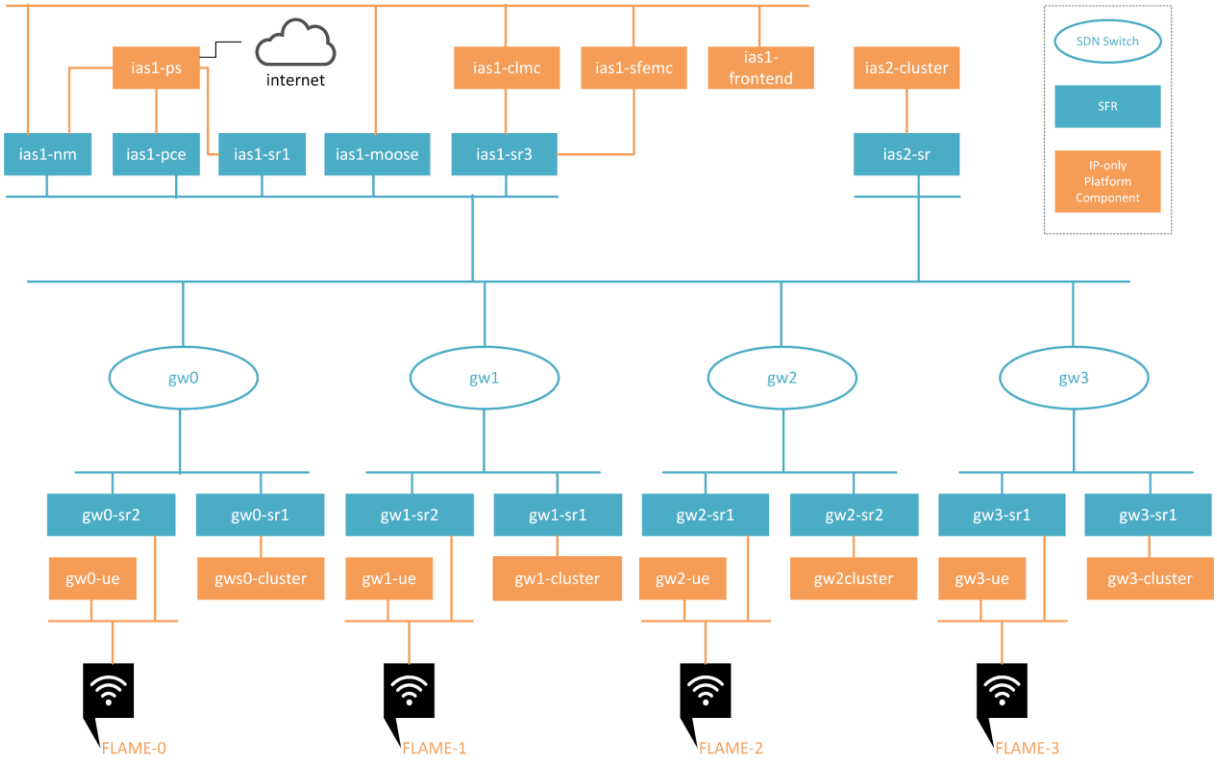

Figure 21: Platform topology in Barcelona.

6.4 LONDON

KCL is one of the two open call two winners. KCL offers four Wi-Fi access points with a dedicated OpenStack compute for each location. Additionally, KCL operates with a hardware SDN switch in their infrastructure interconnecting the edge compute nodes with the data centre one.

6.4.1 Infrastructure Slice

The slice created for the FLAME tenant has an SDN hardware switch at the centre which interconnects each compute node on a dedicated link, as illustrated in Figure 22. Each edge compute node *(Low Latency/Standard Far Edge N*) is serving a particular Wi-Fi access point with the ability to compare the deployment of the platform and SFEs on compute nodes configured with a real-time kernel and standard ones.

Co-funded by the Horizon 2020 Framework Programme of the European Union

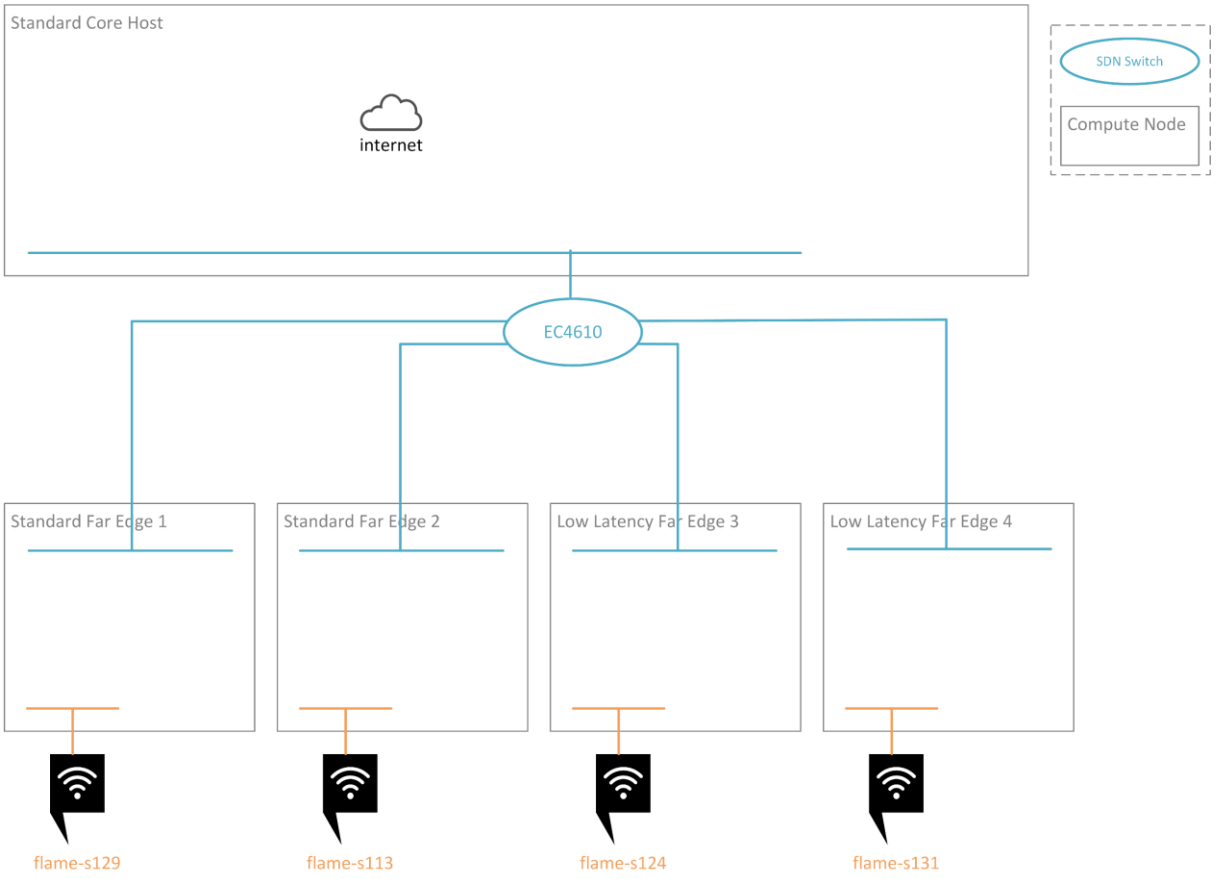

Figure 22: Infrastructure slice for FLAME in London.

6.4.2 Platform

Similar to the deployments in Bristol and Barcelona, each compute node at KCL receives a cluster with the edge compute nodes and emulated UE instance for remote testing. The compute resource for each cluster are given in the table below.

Table 2: Cluster compute resources at KCL

Cluster Name	vCPUs	RAM [GB]	Disk [GB]
dc-cluster	8	22	200
fe1-cluster	5	16	300
fe2-cluster	5	16	300
fe3-cluster	5	16	300
fe4-cluster	5	16	300

Co-funded by the Horizon 2020 Framework Programme of the European Union

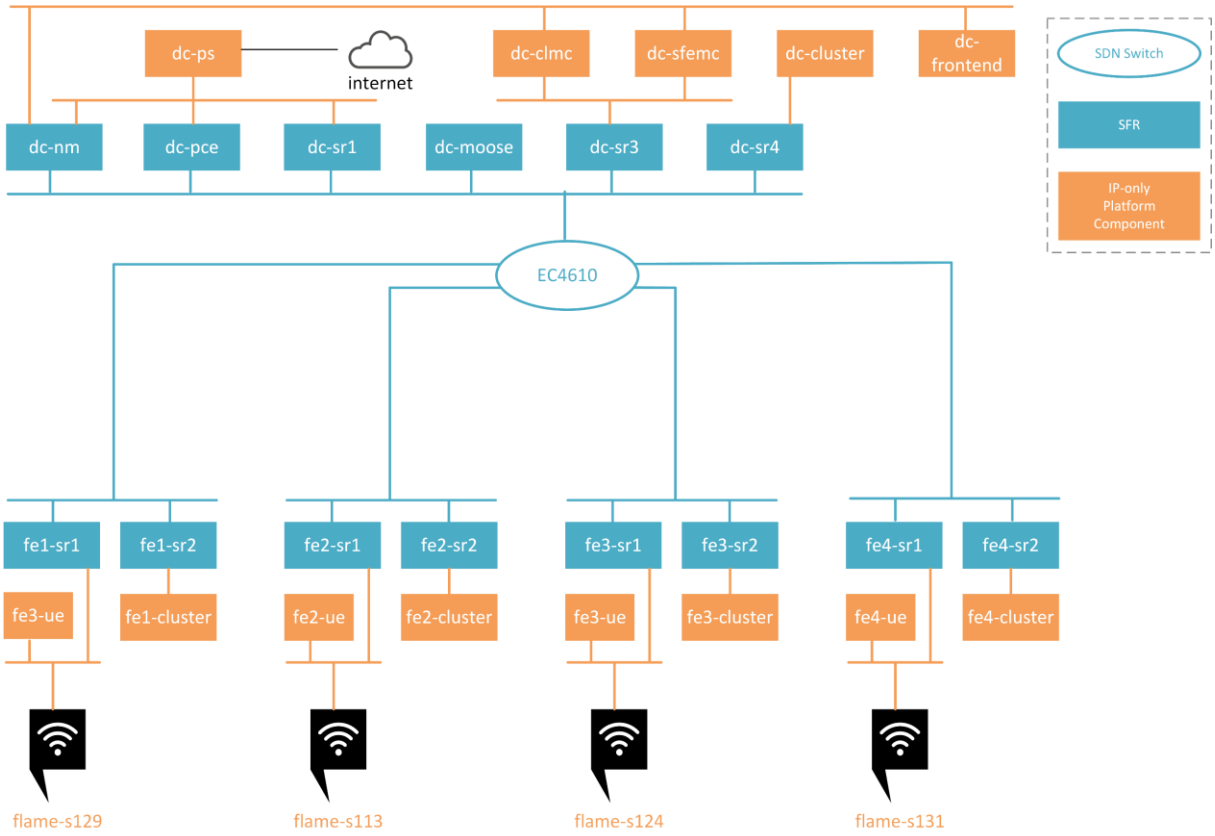

Figure 23: Platform topology in London.

6.4.3 Hardware Installation Timeline

The installation timeline for KCL is given below.

Initiating an Experiment	**Timeline**
Business Proposal	
Experimenter to provide SoW - specifications of infrastructure and use cases. Reviewed by Engineers for viability within deployed infrastructure.	*Completed*
Engineer to provide Solution/Network Design - Include hardware/equipment needed for project and who will provide each element and service	*Completed*
Quote for any contractors involved in the deployment process.	*Completed*
Proposal to include all of the above and be signed off by both parties including any warranties and SLAs. Commencement of project date agreed within proposal.	*Completed*
Delivery Planning	
Site use approval from local government or host building - managed by Engineering where required for technical confirmations	*N/A*
Survey report based on approved deployment sites and any other required surveys	*N/A*
Procurement of the required equipment and resource planning - confirm exact spec and amount of equipment with Engineering	*Completed*
Confirmation of the installation timelines and final costs from suppliers - inform the relevant parties if different from initial quote	*N/A*

Co-funded by the Horizon 2020 Framework Programme of the European Union

Delivery - Installation and Validation	
Installation and testing on site - to include internal architecture/network and confirmation of results (timeline depending on extent of installation)	*First version to be completed. Nov 2019*
Validation against experimenter acceptance criteria outline in initial proposal.	*Jan 2020*
Ongoing support - by engineering team, escalated should changes need to be made to original specs and timeline.	*Jan-May 2020*
Project Closure	
Written confirmation of acceptance and feedback from customer	*June 2020*
Collate lessons report and feedback from customer	

6.5 SICILY

Level7 is one of the open call two winners to replicate FLAME. The deployment in Sicily is a rather unique one compared to all other sites as it spreads over the island of Sicily with the Level7 headquarters in Palermo where the OpenStack data centre compute node is located (also a local node in Palermo is available for testing) and a large installation in Buseto Palizzolo, a rural community close to Palermo

Each location in Buseto Palizzolo comes with one or more own Wi-Fi access point(s) and a compute node with KVM as the hypervisor for the platform instances.

6.5.1 Infrastructure Slice

At Level7's headquarters OpenStack orchestrated compute nodes are available (as well as a local node) with the remaining infrastructure nodes deployed out in the field in or near Buseto Palizzolo. Each node received the name from its location it is deployed and all links between them are Wi-Fi-based connections. As can be seen in Figure 24, the cemetery SDN software switch connects the three locations school, library and municipality over the same link because they are inter connected via a single Wi-Fi carrier. Only from the municipality the square node can be reached.

Co-funded by the Horizon 2020 Framework Programme of the European Union

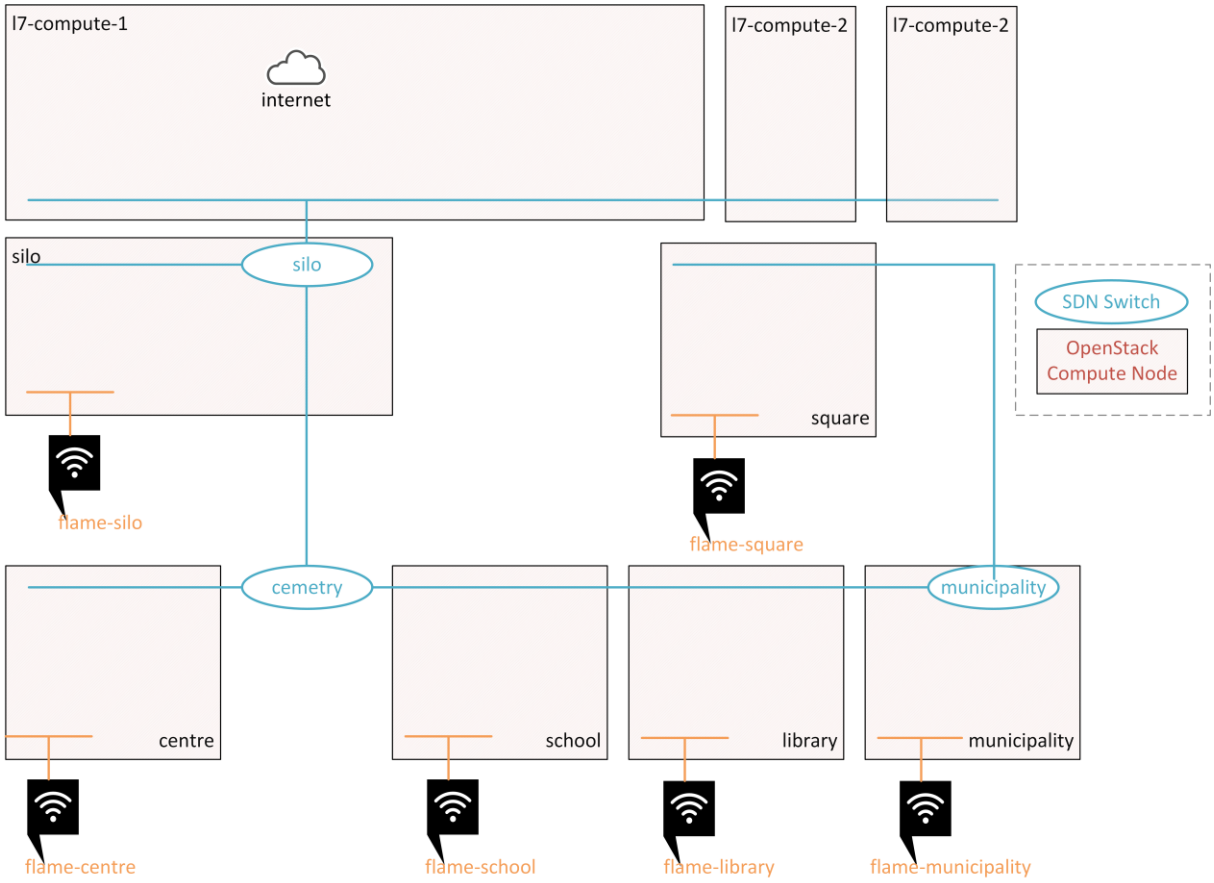

Figure 24: Infrastructure slice for FLAME in Sicily.

6.5.2 Platform

The nodes in Palermo and Buseto Palizzolo are OpenStack orchestrated, therefore methodology to deploy a cluster at each access point location and an emulated UE has not changed, as illustrated in Figure 25. The table below lists the compute resources available to experimenters in Sicily.

Table 3: Cluster compute resources in Sicily

Cluster Name	vCPUs	RAM [GB]	Disk [GB]
l7-cluster	56	314	17000
si-cluster	8	16	450
ce-cluster	8	16	450
sc-cluster	24	32	450
li-cluster	24	32	1000

Co-funded by the Horizon 2020
Framework Programme of the European Union

Cluster Name	vCPUs	RAM [GB]	Disk [GB]
mu-cluster	24	32	450
sq-cluster	24	32	450

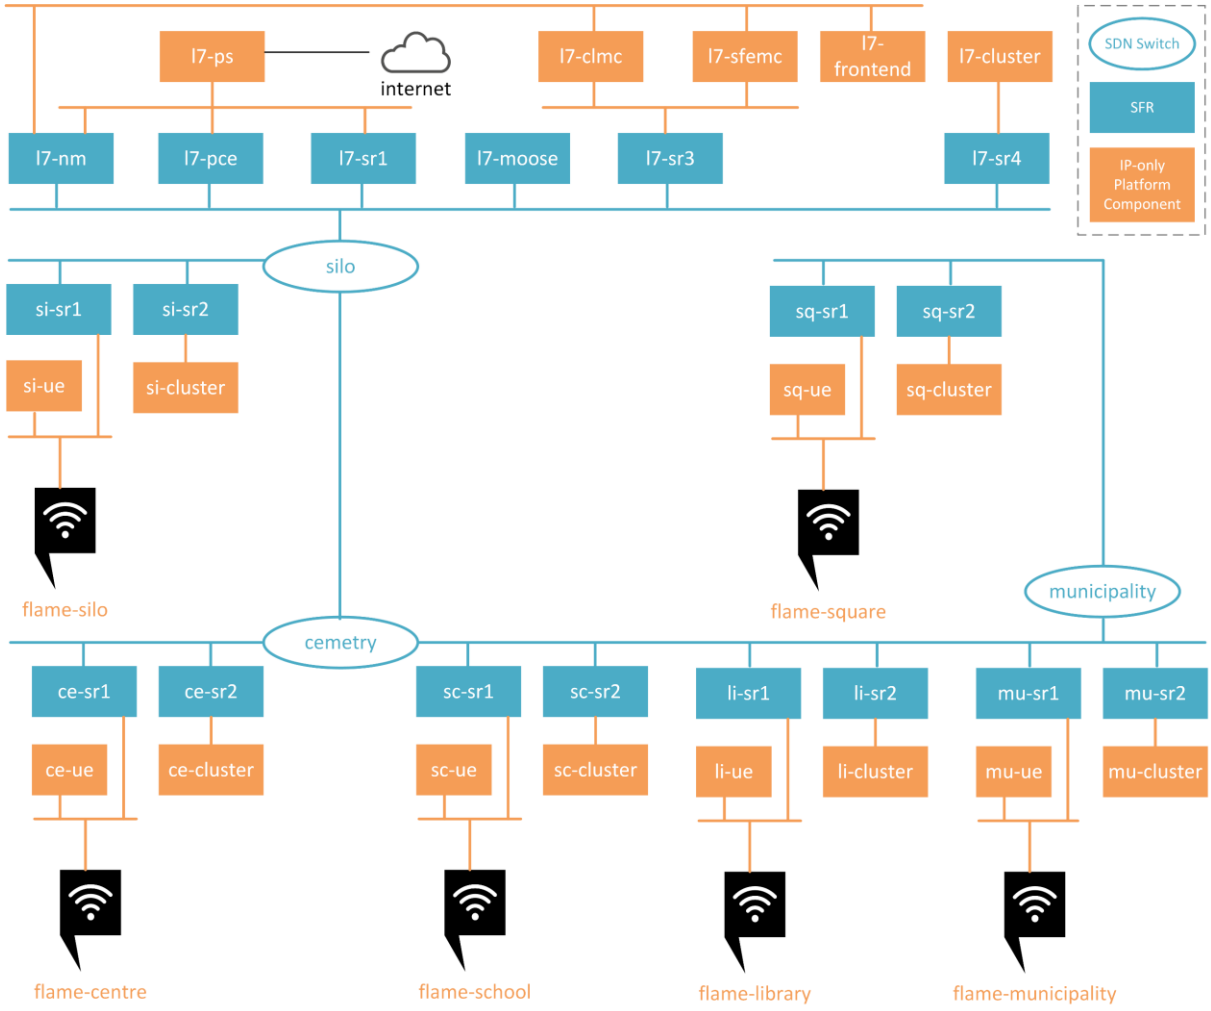

Figure 25: Platform topology in Sicily.

6.5.3 Hardware Installation Timeline

Initiating an Experiment	Timeline
Business Proposal	
Experimenter to provide SoW - specifications of infrastructure and use cases. Reviewed by Engineers for viability within deployed infrastructure.	
Engineer to provide Solution/Network Design - Include hardware/equipment needed for project and who will provide each element and service	*Completed*
Quote for any contractors involved in the deployment process.	*Completed*

Co-funded by the Horizon 2020 Framework Programme of the European Union

Proposal to include all of the above and be signed off by both parties including any warranties and SLAs. Commencement of project date agreed within proposal.	*Completed*
Delivery Planning	
Site use approval from local government or host building - managed by Engineering where required for technical confirmations	*Not Required*
Survey report based on approved deployment sites and any other required surveys	*Not Required*
Procurement of the required equipment and resource planning - confirm exact spec and amount of equipment with Engineering	*Completed*
Confirmation of the installation timelines and final costs from suppliers - inform the relevant parties if different from initial quote	*Completed*
Delivery - Installation and Validation	
Installation and testing on site - to include internal architecture/network and confirmation of results (timeline depending on extent of installation)	*Data Center: (Completed) Nodes: November 2019 (In progress)*
Validation against experimenter acceptance criteria outline in initial proposal.	*Q2 2020*
Ongoing support - by engineering team, escalated should changes need to be made to original specs and timeline.	*Jan-Sept 2020*
Project Closure	
Written confirmation of acceptance and feedback from customer	*September 2020*
Collate lessons report and feedback from customer	

Co-funded by the Horizon 2020 Framework Programme of the European Union

7 CONCLUSIONS

This document has been designed in a way that replicators can easily understand what the requirements of FLAME are for an infrastructure and which steps have to be followed to make FLAME operational. Section 2 describes how the SDN switching fabric has to be set up, how compute nodes need to be configured, the requirements of radio access nodes and how OpenStack is used to deploy FLAME in an infrastructure. With these requirements stated, a series of recommendations are made on how to analyse a new site for suitability in the replication process. The description of how the requirements are met is given for each of the active FLAME replicators, followed by a description of tests and methods to test the operational readiness of a deployment.

Once an infrastructure has been enabled to support FLAME, a deployment workflow needs to be followed, implemented by the ARDENT tool. The workflow and usage of the new version of ARDENT are described in Section 3, offering new replicators an easy-to-use interface to deploy FLAME.

Section 4 presents a series of valuable insights obtained from the deployment of FLAME in the four FLAME infrastructures. These insights reveal that certain aspects of the replication process, mainly related to the infrastructure, can have a relevant impact on the deployment and some of them might be worthwhile considering when planning a FLAME replication. These aspects can range – among other things - from the design of hardware components, to the choice of where the hardware is deployed to more general topics, such as how public spaces are used for experimentation and how public procurement can impact the timeline of a FLAME replication process.

Assuring the experimenters readiness to perform experiments in a FLAME deployment is another key aspect of the replication process that is covered in Section 5. Over the course of the project, an elaborated workflow has been established that guides experimenters from the initial planning phase of their experiment to the final deployment of the experiment in a replicator site. The intermediate steps to get there are described in detail, from the use of FLAME-in-a-Box, over the Sandpit and the use of emulated UEs in the actual infrastructure.

Finally, Section 6 gives an overview of the replication progress of each city in which FLAME has been deployed so far, including the updated timelines for the new replicators in London and Sicily (Buseto Palizzolo).

Co-funded by the Horizon 2020
Framework Programme of the European Union

8 REFERENCES

[FLAME-D3.10] D3.10: FLAME Platform Architecture and Infrastructure Specification v2, available at https://www.ict-flame.eu/download/d3-10-flame-platform-architecture-and-infrastructure-specification/?wpdmdl=1515&masterkey=5c790eda75dbb

[FLAME-D3.11] D3.11: FLAME Platform Architecture and Infrastructure Specification v3, available at https://www.ict-flame.eu/download/d3-11-flame-technology-roadmap-v1-1/?wpdmdl=2019&masterkey=5dbf43b1cffa2

[FLAME-D5.1] D5.1: FLAME Replication Process v1, available at https://www.ict-flame.eu/download/d5-1-flame-replication-process/?wpdmdl=907&masterkey=5a89b90000bd5

9 APPENDIX

9.1 ARDENT DESCR

9.1.1 ardent descr add

Status ID	Status String
0	Request Successful
1	Payload is expected in request
11	Failed to read bytes from infra-descriptor file
101	Request is already in progress
112	Failed to add descriptor to DB
151	compute_nodes is empty in infra descriptor
151	Failed to unmarshal infra descriptor

9.1.2 ardent descr delete

Status ID	Status String
0	Request Successful
1	Payload is not expected in request
101	Request is already in progress
101	HEAT stack is launched
121	Failed to list Stack
121	tenant-openrc file is not uploaded
121	Failed to unmarshal OpenStack command output

Co-funded by the Horizon 2020 Framework Programme of the European Union

9.2 ARDENT RC

9.2.1 ardent rc admin|tenant add

Status ID	Status String
0	Request Successful
101	Request is already in progress

9.2.2 ardent rc admin|tenant delete

Status ID	Status String
0	Request Successful
1	Payload is not expected in request
101	Payload is not expected in request

9.3 ARDENT CHECK

9.3.1 ardent check run

Status ID	Status String
1	Payload is not expected in request
101	Request is already in progress
101	Incomplete infra descriptor found: compute-nodes are not present
401	Incomplete infra descriptor found: networks is empty
411	Infra Context not initialized properly

9.3.1 ardent check status

Status ID	Status String
0	Check finished

Co-funded by the Horizon 2020
Framework Programme of the European Union

411	Error in reading sanity result file
411	Error in reading sanity result file
451	Sanity-Check failed

9.3.2 ardent check results

Status ID	Status String
0	Results fetched successfully
1	Payload is not expected in request
401	Sanity-Check not initiated
401	Sanity-Check is in progress
411	Error in reading sanity result file
411	Error in unmarshalling sanity result
451	tenant-openrc file is not uploaded

9.4 ARDENT HOT

9.4.1 ardent hot generate

Status ID	Status String
0	Request Successful
1	Payload not expected in request
11	Failed to write bytes into HEAT template file
12	Failed to add cluster flavors to DB
12	Failed to delete cluster flavors from DB
12	Failed to add node password to DB

Co-funded by the Horizon 2020 Framework Programme of the European Union

301	Incomplete infra descriptor found: compute-nodes are not present
301	Sanity-Check results has not passed successfully
301	Request is already in progress
301	Sanity-Check results has warnings
301	Error in unmarshalling sanity result

9.4.2 ardent hot show

Status ID	Status String
1	Payload not expected in the Request
301	Failed to locate HEAT template

9.4.3 ardent hot delete

Status ID	Status String
0	Request Successful
1	Payload not expected in the Request
301	Failed to locate HEAT template
301	HEAT Stack has already been launched
321	Failed to list Stack

9.5 ARDENT STACK

9.5.1 ardent stack create

Status ID	Status String

Co-funded by the Horizon 2020 Framework Programme of the European Union

1	Payload is expected in the request
11	Failed to decode json received in request body
201	Request is already in progress
201	Request is already in progress
201	HEAT template is not generated
201	Stack already exists
212	Failed to retrieve Flavors from Storage
221	admin-openrc file is not uploaded
221	Failed to list OpenStack Flavor
221	Failed to unmarshal OpenStack command output
221	Failed to create flavour
221	Failed to create stack

9.5.2 ardent stack delete

Status ID	Status String
1	Payload is expected in the request
11	Failed to decode json received in request body
201	Request is already in progress
201	Empty stack name
201	Stack does not exist
221	tenant-openrc file is not uploaded
221	admin-openrc file is not uploaded

Co-funded by the Horizon 2020 Framework Programme of the European Union

221	Failed to list openStack Flavor
221	Error in listing stack
221	Failed to delete Stack
221	Failed to delete flavor

9.5.3 ardent stack status

Status ID	Status String
0	<STATUS>
1	Payload is not expected in the request
201	Stack does not exist
221	tenant-openrc file is not uploaded
221	Failed to list Stack
221	Failed to unmarshal OpenStack command output
221	Failed to get stack status

9.6 CHANGELOG

This subsection lists the changes made compared to the draft version submitted in November 2019. In more general terms the document has undergone through minor editing (spell and grammar checks) that are not denoted in the following table.

Section	Changes Made
Section 1: Introduction	Minor editing to reflect the updated content of the deliverable (also in the Executive Summary)
Section 2: Enabling an Infrastructure for FLAME	Details about SDN-fabric added for Sicily and London Infrastructure details added for Sicily and London

Co-funded by the Horizon 2020 Framework Programme of the European Union

	Operation readiness explained for Sicily and London Format corrections (minor)
Section 3: Platform Deployment	No updates in this Section.
Section 4: Specific Deployment Insights	Lessons learned in Sicily and London during the operation of the FLAME platform on top of their infrastructures have been added.
Section 5: Assuring Experimenters' Readiness	No updates in this Section, as it was already finished for the draft version.
Section 6: Replicator Progress	The replication progress has been updated for Sicily and London to reflect the latest status.
Section 7: Conclusions	Updated the conclusions Section to reflect the final status of the deliverable content
Section 8: References	Updated reference to D3.11 with the now available link to the PDF.
Section 9: Appendix	Added a changelog section (9.6)

Co-funded by the Horizon 2020
Framework Programme of the European Union