



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



FLAME Replication Process v1

August Betzler, Pouria Sayyad Khodashenas, Carolina Fernandez, Marisa Catalan (i2CAT)
Gonzalo Cabezas, Mariano Lamarca (IMI)
Katherine Rowbotham, Isaac Fraser, David Jones (BRISTOLOPEN)
Sebastian Robitzsch (IDE)

29/01/2018

This deliverable provides the first version of the FLAME Replication process document. This document aims to present guidelines and recommendations for the replication of the FLAME platform in future Smart Cities. This initial version focuses on the experiences of Bristol (as the FLAME departure city) and Barcelona (as the first replicator). In addition to the general recommendations, some initial guidelines about the FLAME implementation in Bristol and Barcelona are provided as an example for future replicator cities. This document is directly linked with the work performed in WP4 (FLAME platform development and implementation) and includes some of the initial outcomes of WP2 about sustainability, governance and exploitation.

Work package	WP 5
Task	T5.4
Due date	31/12/2018
Submission date	29/01/2018
Deliverable lead	I2CAT
Version	1.0
Authors	August Betzler, Pouria Sayyad Khodashenas, Carolina Fernandez, Marisa Catalan (i2CAT), Gonzalo Cabezas, Mariano Lamarca (IMI), Katherine Rowbotham, Isaac Fraser, David Jones (BRISTOLOPEN), Sebastian Robitzsch (IDE)
Reviewers	Tomas Aliaga (Martel), Steven Poulakos (DRZ)
Keywords	Implementation, Infrastructure, Replication, SDNs, requirements, guideline, platform

DISCLAIMER

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

This document reflects only the authors' views and the Commission is not responsible for any use that may be made of the information it contains.

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	

EXECUTIVE SUMMARY

This document provides a set of best practices, technical descriptions and guidelines that aim to ease replication of a FLAME-capable infrastructure. The descriptions and guidelines are based on the current status of a production environments (Bristol) as well as integration plans, preliminary testing and deployment results for the Barcelona infrastructure. Further, suggestions and recommendations also aim to reach beyond the scope of the Bristol and Barcelona deployments, e.g. even larger deployments and the use of other communication technologies. The deliverable also gives insights on initial evaluations and performance measurements carried out in the Bristol and Barcelona deployments, as the experience gathered in these evaluations might serve as orientation to parties interested in replicating the FLAME architecture in their infrastructure.



TABLE OF CONTENTS

1 INTRODUCTION..... 10

1 REPLICATION PROCESS – BUSINESS PERSPECTIVE..... 11

2 REPLICATION PROCESS – TECHNICAL PERSPECTIVE..... 24

3 DEPLOYMENT AND COMMISSIONING 28

4 RECOMMENDED SOFTWARE AND TOOLS..... 30

5 INFRASTRUCTURE IMPLEMENTATION EXPERIENCES..... 32

6 CONCLUSIONS AND RECOMMENDATION 52

7 REFERENCES..... 53

LIST OF FIGURES

FIGURE 1: BARCELONA MAP SHOWING THE STREET SEGMENT COVERED IN THE LOW RESOURCE SCENARIO.....	13
FIGURE 2: BRISTOLOPEN FLAME SCENARIO.	15
FIGURE 3: DETAIL OF THE CABINET. BRISTOLOPEN TESTBED.	33
FIGURE 4: BRISTOL HIGH LEVEL INFRASTRUCTURE	34
FIGURE 5: PLATFORM INTEGRATION AT BRISTOL	35
FIGURE 6: BARCELONA FLAME ON STREET DEPLOYMENT AT PERE IV STREET	37
FIGURE 7: VIEW INSIDE THE BOX CONTAINING THE WIRELESS NODES DEPLOYED IN BARCELONA.....	38
FIGURE 8: CABINET SERVER.....	39
FIGURE 9: CISCO ASR920 SERIES	39
FIGURE 10: DC SERVER (SAMPLE)	40
FIGURE 11: DC VIRTUALISATION CLUSTER NODES	40
FIGURE 12: DEPLOYMENT OF THE OPENSTACK COMPONENTS IN THE BARCELONA MAIN DC.....	41
FIGURE 13: INTERFACE AND BRIDGES BONDING ACROSS THE OPENSTACK SERVERS	42
FIGURE 14: CONNECTIVITY BETWEEN SERVERS IN THE OPENSTACK CLUSTER	43
FIGURE 15: HIGH-LEVEL VIEW OF THE FLAME INFRASTRUCTURE IN BARCELONA.....	44
FIGURE 16: NETWORK DIAGRAM OF THE MAIN DC IN BARCELONA.....	45
FIGURE 17: NETWORK CONNECTIONS AND NETWORKING EQUIPMENT	46
FIGURE 18: FLAME NETWORKING AND COMPONENT DETAILS FOR THE BARCELONA DEPLOYMENT	47
FIGURE 19: SETUP OF THE WIRELESS NODES FOR THE PERFORMANCE TEST IN PERE IV.	48
FIGURE 20: RESULTS FOR THE 40 M BACKHAUL LINK	49
FIGURE 21: RESULTS FOR THE 107 METERS BACKHAUL LINK	50
FIGURE 22: RESULTS FOR THE 205 M BACKHAUL LINK	51



LIST OF TABLES

TABLE 1. ESTIMATED TIMELINE FOR A MEDIUM RESOURCES SCENARIO.....19

ABBREVIATIONS

3PPM	3rd Party Project Manager
APN	Access Point Name
BCC	Bristol City Council
CAPEX	Capital Expenditure
CoAP	Constrained Application Protocol
CPU	Central Processing Unit
DNS	Domain Name System
DPA	Data Protection Authority
DC	Data Centre
DVR	Distributed Virtual Routing
EC	European Commission or European Community
EaaS	Experimentation as a Service
EEA	European Economic Area
EM	Ethics Manager
EMB	Ethics Management Board
EPC	Evolved Packet Core
EU	European Union
ETSI	European Telecommunications Standards Institute
FDD	Frequency-Division Duplex
FLIPS	Flexible IP-Services
FMI	Future Media Internet
GPS	Global Positioning System
GW	Gateway
ICO	Information Commissioner's Office
IP	Internet Protocol

IoT	Internet of Things
KPI	Key Performance Indicator
KVM	Kernel-based Virtual Machine
LAN	Local Area Network
LWM2M	Lightweight Machine-to-Machine
LiFi	Light Fidelity
LoRa	Long Range Technology
LoS	Line of Sight
LTE	Long Term Evolution
MEC	Multi-Access Edge Computing
NAP	Network Access Point
NFV	Network Function Virtualization
ODL	OpenDayLight
OPEX	Operative Expenditure
OSM	Open Source Mano
PAN	Personal Area Network
PC	Project Coordinator
PDN	Public Data Networks
PIA	Privacy Impact Assessment
PIML	Personalized, Interactive, Mobile and Localized
PM	Project Manager
PMB	Project Management Board
RAN	Radio Access Network
SBC	Single Board Computer
SFP	Small Form-Factor Plugabble
SME	Small and medium-sized enterprises
SDN	Software Defined Networking

SSID	Service Set Identifier
TCAM	Ternary Content-Addressable Memory
UTP	Unshielded Twisted Pair
VLAN	Virtual Local Area Network
VM	Virtual Machine
VNF	Virtual Network Function
VPN	Virtual Private Network
Wi-Fi	Wireless Fidelity
WP	Work-package
WPL	Work-package Leader

1 INTRODUCTION

The purpose of this document is to provide initial guidelines and recommendations that can be followed by potential replicators of the FLAME platform. To do so, the document will describe the deployment and replication work performed in Bristol and Barcelona as two different examples on how a city infrastructure can be integrated with the FLAME platform. Note that at the time of this writing, the alpha FLAME platform is in the final stage of development (the release is planned for February 2018) and, thus, Bristol and Barcelona are still working on the integration of the infrastructure. Furthermore, the installation of the city deployment in Barcelona is planned to happen in the first half of 2018. Thus, the following document includes the initial guidelines for replication, technical details about the current integration plans in both cities and some preliminary testing and deployment results. This deliverable focusses on the city of Bristol as the FLAME city infrastructure initiator. The second version of the document, at the end of the project, will include updated details of the experiences in Bristol, the replication in Barcelona and also the insights, expertise and learnings of the rest of replicator cities.

1 REPLICATION PROCESS – BUSINESS PERSPECTIVE

The intention of Section 2 is to provide some general recommendations on the replication process from a business point of view, without going into the details of the technical implementation. These recommendations cover different aspects, such as regulations, timelines, financing and skill requirements that should be taken into account by interested cities or infrastructure deployments before replication.

2.1 MINIMUM DEPLOYMENT EXPECTATIONS

It is expected that the initial process taken on by replicators will focus on deployment of hardware within the physical city setting and the implementation of virtual resources. Bristol Is Open (BRISTOLOPEN) has specified the technical and operational factors that they consider in the deployment process for experimentations within their test-bed infrastructure. It is expected that similar considerations will need to be made when replicating the FLAME platform within other city settings. These include a variety of operational and technical considerations and the interaction between physical and technical challenges.

Cities joining the FLAME project as replicators at the very least need to consider the factors below to be able to offer Experimentation as a Service (EaaS) and sustain the FLAME platform beyond the initial project.

Stakeholder Participation

FLAME replication cities will need to identify those stakeholders who are integral within their city for the implementation of the FLAME platform. It is expected that the main stakeholders will be the local citizens and government. Identifying the benefits of integrating the FLAME platform for these stakeholders will be crucial in the process of deploying the platform within the replication city.

Within the two test-beds of FLAME, the ownership of the city assets (lampposts and street cabinets) where hardware is deployed belongs to the local governments. The expectation in replication cities is that approval will need to be sought from the local governments for hardware to be deployed on these assets. This will need to be coordinated with the local governments, i.e. approval for installation of nodes on lampposts or street cabinets with city maintenance managers or utilising their preferred contractors for installation. The local government will need to take into consideration the health and safety, public right of way, and cost of repairs. These factors may inhibit timelines of experimentation and should be considered in project planning.

Within Bristol, the test-bed infrastructure also relies on the collaboration of local stakeholders, who host hardware; this includes the location of optical fibre access points at the science museum or business incubation sites. These relationships are crucial for maintaining the optical and Wi-Fi access points, and allowing engineers to troubleshoot issues with hardware.

Financing

A financial plan should be an initial consideration for FLAME replication to ensure self-financing beyond EU funding.

Analysing the initial financial cost of replicating the FLAME platform within a city should identify the viability of integration and aid the procurement of crucial stakeholder interested in the project.

Identifying the initial set-up costs and running an analysis of the costs of sustaining the platform will identify the profitability of the platform to intended stakeholders against intended outputs for end-users.

To deploy the hardware and software initially required for the FLAME platform integration within a smart city test-bed environment will be a large upfront cost. Considerations should be made by replicators to the model of public private partnerships, as commercial technology companies may help in the provision of hardware and software, especially as the platform evolves over time and expands based on user demand.

Regulatory and Admin

Deployment of the FLAME platform will need to accommodate within their timeline of deployment their local and national regulatory guidelines. It is expected that the various replicator cities may have different governing stakeholders requiring different forms of regulatory and compliance from the FLAME replicator. We propose that replicator cities will need to consider the factors below on a national and local level when planning their deployment. These should be incorporated into the deployment timeline as soon as possible, as they may pose a challenge to the swift deployment of the FLAME platform.

- **Ethics & Security** – Ensuring public safety in terms of physical environment and personal data security. These regulations will be put in place by the local governing body to protect the rights of their citizens. Ethical and security clearance will be needed if deploying cameras or microphones in public places and the deployment of hardware to be hosted in public spaces, i.e. active nodes.
- **Technical Hardware Licensing** – In the UK, an RF license is required to deploy hardware that emits a radio frequency. The process of procuring such a license may vary across EU countries and cities but it is expected that this will need to be in place prior to the deployment of hardware.
- **Health and Safety** – As hardware devices are likely to be deployed in public spaces it is expected that local governing body or commercial owners of these spaces will require the installer to comply with national and regional health and safety legislation. It is possible for the replicator city to leverage this process by contracting an installation specialist for hardware whom the local governing body are preferential to, as they will have health and safety training and understand the processes of the local government to ensure public safety.

Health and safety legislation which would impede engineers to deploy hardware and access to lampposts themselves may require machinery. Bristol, to efficiently and safely deploy hardware, employs the local government's contractors to carry this out. This should be considered for replication cities that may not have the skills required to install hardware at dangerous height or operate electrical equipment on the street. This option also provides minimal disruption to local citizens and considers methods to protect the public from installation.

2.2 RECOMMENDATION OF DIFFERENT SCENARIOS

Three types of scenarios can be distinguished depending on the capacities or willingness of replicators to invest on equipment and infrastructure. We refer to these as low, medium or high resources scenarios. As stated in [1] (Section 2.3.1), two kinds of infrastructure can be considered for FLAME: 1) a software-based, which offers high flexibility and lower cost, and 2) hardware-based, which requires dedicated hardware resources but can offer a higher performance. The software-based kinds of

infrastructure might be considered for a low resources scenario (e.g. for a small testbed deployment in cities that do not depart from an already deployed infrastructure or that cannot afford notable hardware investments). The hardware-based infrastructures might be more appropriate for medium or high-resources scenarios. Herein, the three types of scenarios are described.

Low resources scenario

The low resources scenario is based on the test-bed environment of i2CAT located on a street in the district of Sant Martí in Barcelona, Spain.



Figure 1: Barcelona map showing the street segment covered in the low resource scenario

Barcelona is located in the North-East of Spain and counts around 1.6 million habitants. The district of Sant Martí is the 4th largest district. The district has a high population density and the test-bed deployment zone is an urban area with small and medium commercial buildings. The test-bed stretches along approximately 500 meters of the Pere IV street (Figure 1).

The testbed features the following technical capabilities:

- Wireless connectivity – a wireless network that is built with 5 Wi-Fi nodes that provide Radio Access Network (RAN) connectivity and are connected to each other over a wireless backhaul.
- Fibre connectivity – fibre connectivity between the main Data Centre (DC), the edge resources and to each Wi-Fi node in particular is provided.
- Switching technologies – switching nodes are installed to provide connectivity between the street site and the resources in main DC resources.
- Server-based resources – 3 dedicated servers in the main DC and one edge server

- Network slicing capabilities – network slicing to separate application (verticals), management and service aspects of the network is applied via virtual LANs (VLANs).
- Network Visualization – tools like Grafana [2] will be used to provide a platform for monitoring and analytics.

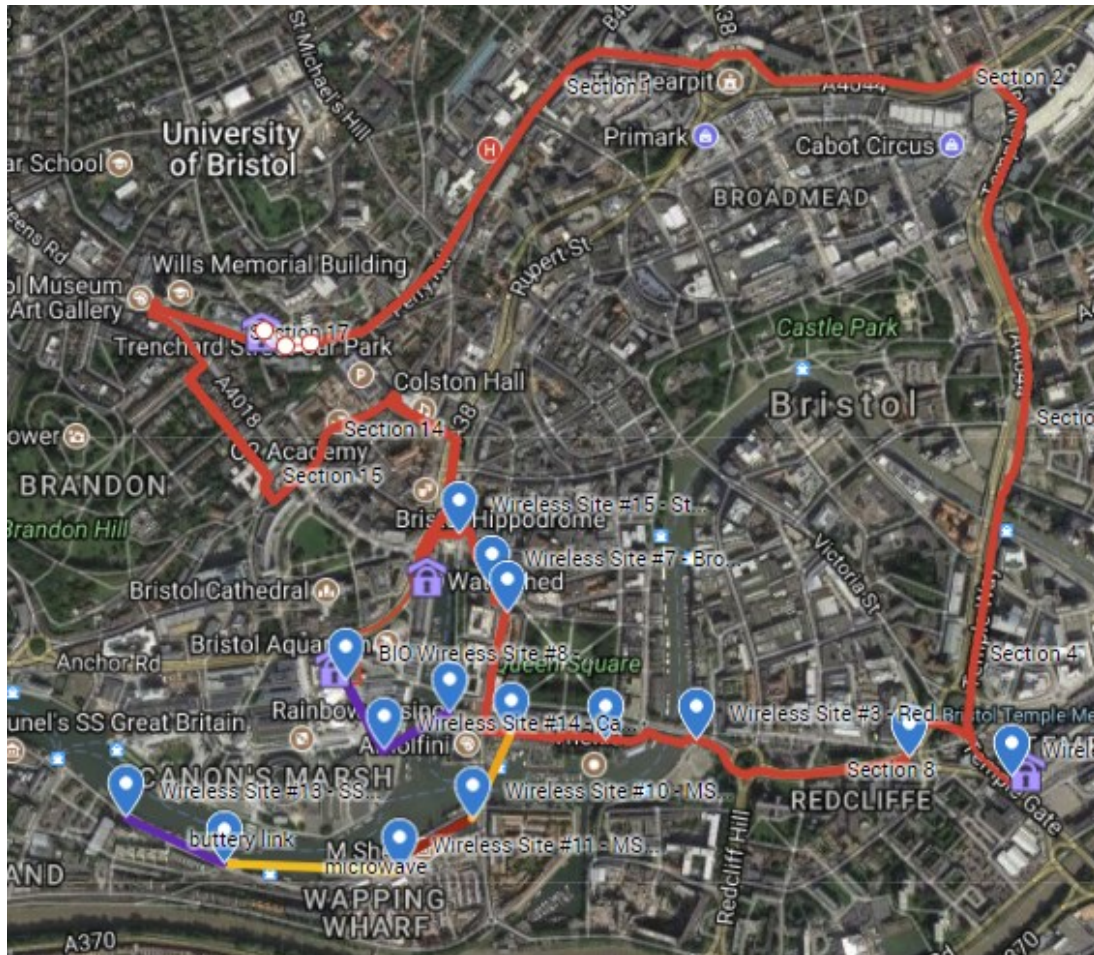
To provide the features listed above, the following resources are used:

- Wireless Nodes: Wi-Fi nodes used both for the RAN and the backhaul are deployed on lampposts along the street. The nodes consist of a single board computer (SBC) equipped with two or more Wi-Fi cards supporting the IEEE 802.11ac (and a/g/n) standards. The SBC is contained within a custom weather-proof box that also hosts a media converter from Ethernet to fibre, a remote control module to reset the nodes on demand, as well as battery and power converter units. The nodes on the lampposts are connected via fibre to the cabinet server.
- Cabinet computing element: To provide on-site computing resources, a server is placed in a street cabinet on one of the crossing of Pere IV. The server is connected via Ethernet to the rest of the network and it provides computing resources that are available to OpenStack [3] . Services that require close proximity to the wireless nodes on the lampposts (e.g. for small delays) can benefit from running on the same street where the lampposts are located.
- Main DC computing elements: In the main DC, servers host the network services, OpenStack, and other computing intensive services like media services. The main DC is located close to the i2CAT premises in a dedicated space (e.g. a server room).
- Networking components: A series of switches and routers are used to provide connectivity between the lampposts, the edge cabinet (including the server) and the main DC resources located at i2CAT. Fibre is used to connect the lampposts with the cabinet and to connect the cabinet to the main DC.

Medium resources scenario

The medium resource scenario is based on the test-bed environment of BRISTOLOPEN.

Bristol the largest city in the south west of England with a population of approximately 450.000 people. The current test-bed environment is deployed to a predominantly urban area occupied by commercial and council-owned buildings. Figure 2 shows the map of the area.



Key

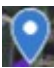
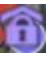
Wi-Fi Active Nodes	
Switching Node	

Figure 2: BRISTOLOPEN FLAME scenario.

BRISTOLOPEN currently offers the following technical capabilities within this environment:

- Wireless connectivity – wireless network hubs located across multiple locations within Bristol City Centre through cell radio technology (see locations above).
- Fibre connectivity – fibre ring connecting to the Wi-Fi active nodes.

- Switching technologies – four switching nodes are deployed within areas of the city. Server based resources – utilises OpenStack as a cloud controller provider environment for experimentation across the BRISTOOPEN platform and hardware.
- Network slicing capabilities – the BRISTOOPEN platform through OpenStack has slicing capabilities using VLAN's.
- Network Visualization – BRISTOOPEN has several options to provide monitoring and network visualisation across the platform infrastructure to experimenters to allow for management of experimentation and troubleshooting and diagnosis.

To host the above technologies, BRISTOOPEN has several physical resources to host and deploy the hardware. This is based around a network of multiple points of presence across the area above.

- Wi-Fi / Council Assets – The devices facilitating these capabilities are hosted on street lamp posts and street cabinets. These resources are made available by the council and are installed by a private contractor. Within the project planning stage, BRISTOOPEN will need to work with the council and installation company to gain approval to use and operate them. BRISTOOPEN will secure the street cabinets and will have permission to host hardware within the street cabinets from the council. Access and functionality at these sites may be dependent on ongoing street level maintenance on traffic systems. To ensure the maintenance of the network, engineers will need to coordinate with the council to plan for maintenance works. BRISTOOPEN is limited to certain areas of the city depending on coverage and access requirements based on the council's agreements.
- Switching and Active Nodes / Partner Server Rooms – The switching nodes are installed in server rooms of partners of BiO, namely, University of Bristol, Science Museum, Water Shed and Engine Shed. Engineers will need to have access to these server rooms through coordination with their estates or IT teams. These teams are to be consulted should new installations be needed for specific projects. This is often in consultation with local contractors who also have a relationship with these partners.

High resources scenario

We propose a theoretical, high resources scenario for any interested replicator that disposes of even more resources than for the medium resource scenario. Much like the medium budget scenario, the high budget scenario is a larger scale deployment of the FLAME platform.

It is expected that the higher resource scenario will have similar capabilities to the Bristol functionalities, but would cover a greater geographical area requiring more Wi-Fi devices and even cellular networks, fibre connectivity and switching points.

- Wireless connectivity – wireless network hubs across multiple locations within and beyond the city centre utilising radio technologies like Wi-Fi, Long Term Evolution (LTE), etc. with licensed and unlicensed spectrums on multiple points. The high resource scenario would be expected to range across a regional area connecting urban areas and satellite towns. This would require the collaboration with different local governments and the utilisation of their street level assets. Requirements would need to be identified to co-ordinate the compliances for each local government organisation e.g. hardware installation requirements, public safety, and legal compliances.
- Fibre connectivity – fibre connections extending the wider range of Wi-Fi devices. In this scenario, the fibre ring would be extended from the city locations to conurbations of the city that would have service users in need of fibre connectivity i.e. community centres, media hubs

and local colleges. This would allow the higher resource replicators to access the fibre access nodes and connect to the larger network of Wi-Fi/Internet of Things (IoT) devices.

- Switching technologies –switching nodes could be deployed within the areas of the fibre extension, as has been done in the Bristol scenario, allowing for engineers to access for maintenance and improvement works within the community centres, media hubs or colleges.
- Network Slicing and Visualization – Given the extended physical resources, this scenario would require increased capability to accommodate the greater demand of experimenters and users on the network. The platform would still require slicing capabilities using VLAN's, but the operator would need to accommodate the maintenance and operation within their operational resources e.g. staffing, compute, memory and storage capabilities. With this greater capability and demand for experimentation, additional monitoring would need to take place across the FLAME platform.

The high resources scenario will be deploying hardware to one or several geographical locations within a metropolitan area. Depending on the size of the city, this would need to be managed taking into account the city limitations such as physical barriers and local government regulations. It is recommended that these should be managed within clusters. For example, defining clusters based on geographical location (North and South Clusters) or depending on the size or capabilities. This management system would not only be for the maintenance requirements but also for the network monitoring and project management.

2.3 DEPLOYMENT REQUIREMENTS

Typically, deployments will involve a variety of infrastructure mechanisms depending on physical constraints of the city setting. Installations should be planned based on coverage and performance required and this will dictate where equipment should be placed.

Wireless and Radio equipment

Wireless and Radio equipment can be deployed on lampposts. In these locations, equipment needs to be ruggedized for outdoors and there has to be an easy way to access (reset) these devices. Devices that can be reset remotely are recommended. There are power, aesthetic and weight considerations to mounting equipment on a lamppost which are usually defined by the lamppost manufacturers specifications. In Bristol, most radio equipment is mounted in a street cabinet at the base of the lamppost for accessibility and security reasons. In Barcelona, the radio equipment is mounted on top of the lampposts within safely attached boxes. Cabinets need to be locked, waterproof and suitable for the location in which they are deployed. These cabinets are typically IP68 rated.

Wireless equipment can also be installed directly on buildings. These installations are easy from access and maintainability perspective but must be somewhere installed where that is secure. In Bristol, our Small and Medium-sized Enterprises (SME) partners host this wireless equipment on balconies and rooftops, allowing for easy access.

Active nodes / Switching capabilities / Firewalls / Monitoring equipment

This range of equipment should be installed in a secure room within a building. Staff should be able to access equipment for regular planned maintenance but also be able to gain rapid access in case of an emergency. Bristol opts to host this type of equipment within SME partner server rooms. In Barcelona, the equipment is hosted within server rooms that are only accessible by the city or to project partners

with access rights. Replicators should consider the cost of implementing this process and the power drawn.

Fibre / radio infrastructure

Fibre and radio infrastructure should be deployed per local regulations, but will involve civil works and these can be costly. The preferred local government contractor is usually recommended as they will be familiar with the assets across the city.

Smart City Technologies

Smart city technologies can include many different high and low-level sensors and actuators. More concretely, it depends on the use cases for the specific implementation. The minimum requirements for use cases would require Wi-Fi connectivity to the test-bed network and FLAME platform. However, further considerations would be needed to utilise IoT sensors depending on use case requirements proposed by experimenters. This is something that should be addressed by replicators based on their resource availability, and at the proposal stage of experimentation.

2.4 SKILL REQUIREMENTS

To successfully deploy infrastructure within a city setting, a range of technical and operational skills will be needed.

Engineering skills

Engineers employed within replication cities will need to cover a range of technical skills as specified below:

- *Computing Environment Engineer* – This engineering role will require experience and skills for cloud environments, DevOps and network management. The engineer should have the ability to troubleshoot extensive issues; e.g. system resources.
- *IoT Engineer* – In large deployments where IoT elements are involved, this role will require knowledge of IoT specific technologies and protocols (Constrained Application Protocol (CoAP), Lightweight Machine-to-Machine (LWM2M), etc.), and IoT specific platforms (FIWARE, Intel, etc.). Additionally, they would require experience with scripting languages (Ruby, Python, etc.), networking and network management, and SQL databases.
- *Core Networks Engineer* – This role will require knowledge and experience of Software Defined Networking (SDN), Network Function Virtualization (NFV), and Wireless Local Area Network (LAN) / Personal Area Network (PAN). Furthermore, experience and awareness of network security is required, as well as network expertise certification such as Cisco Network Administrator.
- *Wireless and Mobile Engineer* – This role might need knowledge of LTE or Wi-Fi depending on the technology used. Radio planning and network commissioning skills are also important.
- *Virtualization Technologies Engineer* – This role has knowledge of tools used for virtualization of network functions and about the management and orchestration of virtualized environments. Typically, frameworks like OpenStack, Open Source Mano (OSM) [4] or others are used for the management of Virtual Network Functions (VNFs), Virtual Machines (VMs), virtual containers and other virtualized elements.
- *System Administration* – This role has a deep knowledge about software to host cloud systems and / or any radio system. Their task is to administrate these software environments in such a way that that engineers have the necessary access rights and resources assigned, permitting

them to run their software and to store data. Further, the system administration is responsible for enabling connectivity to these systems (e.g. VPNs) so that engineers are able to interact with them.

- *Infrastructure Maintenance* – Any use of infrastructure or hardware requires someone or a team of persons that have deep knowledge about the installations used in the project. The hardware can range from networking and computing hardware in server rooms to street furniture equipped with radio components or any other installation that is used to host or connect project elements.

Operational Skills

Considerations should also be made to operational staff who would work alongside the engineers in the deployment process.

- *Project Management* – Ideally, these roles will have technical background or experience of working on EU projects. They would work alongside the engineering roles to plan the deployment and interact with stakeholders ensuring compliance is followed with local governments, and also focus on the management of stakeholder relations.

2.5 TIMELINES FOR DEPLOYING EXPERIMENTATION INFRASTRUCTURE

Outlined below are the proposed timelines for deploying infrastructure within Bristol City Centre based on BRISTOLOPEN's experience. This is a high-level milestone plan and does not contain the technical/engineering requirements dependant on the use case proposed by experimenters. Please note that this will depend on the scale and topology of the infrastructure and city area where the FLAME platform is to be deployed. In a low or high resources scenario, the time frame would differ due to the varying demand to deploy hardware across the replication cities.

Table 1. Estimated timeline for a medium resources scenario.

Initiating an Experiment	Timeline
Business Proposal	
<i>Experimenter to provide SoW - specifications of infrastructure and use cases. Reviewed by Engineers for viability within deployed infrastructure.</i>	1 wk
<i>Engineer to provide Solution/Network Design - Include hardware/equipment needed for project and who will provide each element and service</i>	1wk
<i>Quote for any contractors involved in the deployment process.</i>	2 wk
<i>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.</i>	3 wk
Delivery Planning	
<i>Site use approval from local government or host building - managed by Engineering where required for technical confirmations</i>	5 wk
<i>Survey report based on approved deployment sites and any other required surveys</i>	5wk
<i>Procurement of the required equipment and resource planning - confirm exact spec and amount of equipment with Engineering</i>	6 wk
<i>Confirmation of the installation timelines and final costs from suppliers - inform the relevant parties if different from initial quote</i>	6 wk

Delivery - Installation and Validation	
<i>Installation and testing on site - to include internal architecture/network and confirmation of results (timeline depending on extent of installation)</i>	8 wk
<i>Validation against experimenter acceptance criteria outline in initial proposal.</i>	9 wk
<i>Ongoing support - by engineering team, escalated should changes need to be made to original specs and timeline.</i>	10 wk
Project Closure	
<i>Written confirmation of acceptance and feedback from customer</i>	10 - 14 wk
<i>Collate lessons report and feedback from customer</i>	

The process for deploying FLAME infrastructure has to follow the standard procedures defined by local regulations. Thus, when public bodies have to subcontract the execution of some tasks, and depending on the nature of the contracts required, different procurement procedures might be applied. Nevertheless, once the corresponding budget is approved to be allocated by the public administration, some common high-level steps can be defined as part this process: 1-Technical specification, 2-Proposals evaluation, 3-Supplier selection, 4-Award stage, 5-Execution.

Typically, the implementation of the whole process might take months from start to end. Therefore, this duration has to be considered when planning city activities that might require public procurement processes.

2.6 LICENSING / CERTIFICATION FOR DEPLOYMENT

There are safety and regulatory requirements for deploying radio equipment that apply in Bristol. Equipment must apply to appropriate standards on the safety compliance statement, which must be provided to the council before any equipment can be installed. To operate LTE radios outdoors in Europe, a license must be obtained (typically renewable every 12 months). This license lists the specific frequencies, antenna heights, RF power and area where the trial will operate. In Bristol, the regulatory authority for LTE is the Office of Communications (Ofcom) and the RF License is Band 7 (2.6G Frequency-Division Duplex (FDD), 15 MHz).

All equipment deployed in cabinets is on an unmetered power supply, so it must have a charge code that allows the council to know the power usage. Additionally, there are rules on the types and sizes of cables that can be installed into street cabinets and lampposts.

The following considerations should be made when deployment takes place;

- Connectivity to lamppost of cabinet - fibre, mmWave, microwave – logistics of implementing connectivity.
- Security – multi locking locks and padlocks to the street cabinet – restricting access and ensuring allowed engineers have access for maintenance.
- Power – from where the power will be drawn from for the devices. In Bristol, most of the power is drawn from lampposts.
- Size – need to assess the size of the device to be deployed and ensure that the lamppost or cabinet in question can host the device and has the space to do so.

- Location – proximity of the device to the lamppost and cabinet. Need to ensure cabling is in the correct locating and engineers can access for maintenance.
- Colour – most devices need to be white if on a lamppost.

When using radio equipment in any public space (indoor and outdoor) or when co-located with any external equipment, there is a set of regulations that need to be taken into account. These regulations determine how the spectrum can be used and the limits for radiation emission from the radio equipment:

- The EU recommendations published in 1999 tried to unify maximum radiation criteria at EU level. Nevertheless, these recommendations are not mandatory for the EU member states and, as a consequence, each infrastructure replicator city should evaluate the local regulations to ensure the technical compliance of the solutions.
- In Spain, the Spanish Government launched in September 2001 the Real Decreto 1066/2001[5] for regulating the radio emissions at the Spanish level and defined the maximum power emissions, aligned with the EU criteria. However, the Spanish local governments had the freedom to modify these limits and reduce them according to their considerations. Thus, the Catalanian government launched in May 2001 its own Decreto 148/2001 [6] that limited the maximum power emissions below the Spanish and EU figures. At present, this regulation is the one affecting the Barcelona site.

2.7 CITY LEVEL DEPENDENCIES

Each local government within a replicator city will have regulations with which replicators will need to comply.

Charge Codes

All equipment deployed in cabinets is on an unmetered power supply, so equipment must have a charge code that allows the council to know the volume of the power. The application process may vary within EU countries. The process for gaining the codes within the UK can be found online.¹

Hardware Specifications

There are rules on the types and sizes of cables and the weight, size and appearance of equipment that can be installed into street cabinets and lampposts. These will need to be agreed upon with the local government to ensure it can be deployed to their assets.

Public Safety

There may be limitations to the street level assets that can be used due to proximity to public walkways, businesses, or physical barriers such as a harbour bridge. These constraints will need to be identified by the local government and taken into consideration when deploying the hardware.

Street Works

¹ <https://www.elaxon.co.uk/operations-settlement/unmetered-supplies/charge-codes-and-switch-regimes/>

Everyone on site has a personal responsibility to behave safely, according to local regulations and have duties to protect citizens from dangers derived from the work activity or the installed equipment. Therefore, all activities must include proper arrangements for design (including planning and risk assessment) and management (including supervision) of the works.

For example, in the case of Barcelona, the entities and those working on their behalf that will carry out local works and activities in the city, must inform and agree to these activities with IMI.

2.8 SUSTAINABILITY AND GOVERNANCE

The governance and sustainability of replication cities will be dependent on the structure of the organisations applying as replicators. The determination of the organisation taking the lead of operating the technical assets within the city setting will decide the type of governance model to be utilised. However, the information provided here will look at a basic concept of governance for replication cities. Sustainability of replication cities should look beyond silos of health, transport, environmental, and social currently established within local governments.

Governance

For replication cities, the role that the local governing body takes determines the governance model, as the city environment is regulated and owned by the local governing body. It is assumed for most EU cities applying as replication cities that the local government will own the city assets required for the FLAME platform deployment and will have an internal or a joint organisation taking on the operation of the test-bed environment. BRISTOLOPEN is an external organisation to Bristol City Council (BCC), but the organisation is a joint venture between BCC and the University of Bristol. Subsequently, these organisations are major stakeholders within the organisation, and they require extensive collaboration with the BCC to deploy hardware and software across the city as well as to approve projects dependant on ethical and legal compliances.

Sustainability

To sustain replication cities beyond the FLAME project, integration with the local government, citizens and environment is critical. Assessing methods to serve the city in improving efficiency and quality of life by the local government and SMEs, can lead to projects that integrate replication test-beds with the development of the city setting. These potential projects may arise with investment from local government to integrate test-bed technology with local services and SME economic development.

The viability of these projects will need to be assessed in line with the following factors; the technical capabilities of each replication city, limiting factors within the city setting, the opportunities local governments are seeking out, and major challenges defined within the city. It is crucial the replication cities ensure that they remain relevant to local governments and SMEs, to re-address arising challenges faced by city governments, and local citizens.

2.9 DATA MANAGEMENT POLICIES

Since the Barcelona infrastructure is under construction, in parallel to the deployment activities, i2CAT will identify potential risks for data exposure, and it will develop/improve the initial recommendations

for a data protection plan given in this subsection in accordance with the European guidelines. As an initial step, data will be collected either by direct or indirect means.

At the infrastructure level, data comprises traffic packets (whose content would be, among others, IP within a private network and payloads containing media content served by a 3rd party using the infrastructure provider's resources). Some extra data (phone number, GPS location of the lamp posts and maybe others as well) may be collected to give access to external experimenters to the infrastructure's private networks.

The aforementioned data can be requested and stored to enable the interaction of the user with the infrastructure. These data are restricted to the phone and other metadata required for each specific experiment to run (like the user's location in the street to identify to which Access Point should it connect to). The type of captured data must be explicitly identified and communicated to the user prior to the registration into the network for a specific experiment; along with the duration of its storage time and contact information so that any user can interact with the infrastructure and/or consortia to exert their right to access, modify or delete their data.

Along with the storage of data, all network traffic data coming from/to a user within a specific experiment is subject to traversing the internal network of a given infrastructure. No data shall be communicated between infrastructures.

The regulations on how data are processed shall be defined at the project and/or the infrastructure level (data controllers). The specific procedures will be laid down by the experiment owners/providers (data processors), which will always be in compliance with the norms provided by the data controllers.

2 REPLICATION PROCESS – TECHNICAL PERSPECTIVE

This section focuses on the technical aspects on replication. Thus, the requirements of the platform and the infrastructure elements (connectivity, edge elements and core) are detailed. Note that some of the technologies mentioned herein go beyond the resources currently provided by BRISTOOPEN or i2CAT in the FLAME testbeds. Nevertheless, these are provided as examples of features and technologies that could be incorporated in future deployments or by future replicators.

2.1 SDN /NFV PLATFORM REQUIREMENTS

2.1.1 SDN Fabric

FLAME is built on a truly SDN-enabled networking fabric [7] and implements a stateless switching solution which requires the switches and controller(s) to be at least OpenFlow 1.3-compatible.

As there is no capability verification alliance for OpenFlow (e.g. Wi-Fi alliance) and the OpenFlow 1.3 features are being considered as “experimental”, it is highly recommended to double-check with the vendor of the fabric if the following two features are supported:

- Switches support arbitrary bitmask matching via semantically overloaded IPv6 fields
- Controller supports handling (read and insert) of arbitrary bitmask matching rules

If the replicator’s infrastructure is based on a software-based switching fabric, it is recommended to use OpenvSwitch, which has not shown any compatibility issues. However, hardware switches implement the actual switch in their ternary content-addressable memory (TCAM) tables, which have an OpenFlow compatible API and only one switch is known to support arbitrary bitmask matching, i.e. PICA8 [8] .

As mentioned before, the chosen SDN controller must a) accept the rules communicated via the REST API and b) insert them into the switches. The following controllers have been successfully tested: Floodlight and OpenDaylight (ODL) [9] . ONOS [10] does not support arbitrary bitmasks yet.

2.2 CONNECTIVITY

Different forms of connectivity should be considered depending on experiment needs. The following should be considered:

- Fibre connectivity – Provides high data bandwidths and is proven technology. Installation is typically static and costly as fibre has to be installed in ducts in the street and might have to run for some distance. Different topologies can be used, such as mesh, point to point, and star topologies (or a combination). Fibre is typically terminated in patch panels and a suitable media converter / small form-factor pluggable (SFP) will need to be used.
- Microwave links – these give point-to-point connections and RAN connectivity. Links should be planned to allow for buildings and topography and will typically provide up to 1 Gbps aggregated (depending on the technology used). Installation cost is usually low, as physical install is only needed at either end.
- Millimetre-wave links – these give point-to-point connections. Links should be planned to allow for buildings and topography and will typically provide up to 4 Gbps aggregated

(depending on the technology used). These devices are typically less robust than microwave links as they are a newer technology. Installation cost is usually low, as physical install is only needed at either end.

- Fibre optic switches – to interface between the backhaul technologies and access technologies, a switching fabric is required. NEC and PICA8 switches have been used in Bristol, but other switches could be used considering the following:
 - Devices should support SDN to allow all devices to mesh together.
 - Devices should support VLANs to allow traffic separation.
 - Devices should be sized according to fibre / electrical split per location. Normally fibre is used in the ring and Ethernet towards edge devices.
 - Required reliability / resilience can be provided by multiple switches or by more expensive switches.
- All Internet connectivity in Bristol currently goes through a central firewall. In the future or in other cities this could be different depending on the network topology.

2.3 ACCESS TECHNOLOGIES

A method for connecting devices into the FLAME platform is required. This can be provided by a variety of technologies. Different forms of access should be considered depending on experiment needs. For FLAME, the following should be considered:

- Wi-Fi access – suitable access points that will allow different network slices to connect with different Service Set Identifier (SSID). Ideally, an SSID has its own unique password that connects to a specific VLAN that is isolated from other slices. Domain Name System (DNS) and Dynamic Host Configuration Protocol (DHCP) will be handled within that network slice rather than by a generic server.
- LTE / wireless access – Different phones should only have access to their slice so only an Access Point Name (APN) that allows specific Public Data Network (PDN) contexts. Ideally, access will be achieved through a private test network with a dedicated virtual Evolved Packet Core (EPC) / EPC cluster that can be interfaced through SDN to other technologies. An alternative proposal would be for devices to access through a commercial cellular network with secure applications that can be connected through the internet to the test network. Typically, devices will either authenticate to a test network or a commercial network, but they will not do both.
- Millimetre Wave – this is included in this section as a potential access technology. Selection of appropriate test devices would need to be considered to use this access technology.
- LORA [11] - This is included in this section as a potential access technology. This technology allows coverage of large areas with multiple devices, each device can only transmit for short periods each day however. Typical applications would be to read a meter where the value does not change very often, and the amount of data sent per device is low.
- Li-Fi – This is an indoor solution based on visible light communications that can be used to provide coverage in specific areas.

For all devices, power and connectivity requirements should be considered for end terminations. It is also useful to have an access device somewhere that allows for device validation prior to installation in a city.

2.4 CORE NETWORK REQUIREMENTS

The FLAME platform requires compute devices to host VNFs. The compute solution you deploy should support integration with orchestration software and SDN controllers. FLAME recommends OSM Mano release three as a potential orchestration tool. OpenStack can be integrated with OSM Mano and SDN controllers. BRISTOLOPEN has recently begun working with OSM Mano release two for the EU projects. Suitability may have to be addressed further down the line during replication.

Bristol hosts an OpenStack Ocata environment configured for self-service networks and Distributed Virtual Routing (DVR) with the Neutron OpenvSwitch agent. A degree of high availability is available for project virtual routers that reside on both the controller/network node and the compute nodes. VRRP is used to provide automatic failover of SNAT services to a backup virtual router running on a different node. Only VMs with floating IP addresses will maintain full network connectivity in the case of a network node failure. No other redundancy exists. In Bristol, there are compute resources in 4 separate locations around the city and VMs can be instantiated on any node as required. In Barcelona there exists a single set of compute resources in the core network. A single OpenStack controller/network node hosts the primary (active) virtual routers with backup virtual routers living on the compute nodes. Virtual routers on the network node perform certain routing functions that virtual routers residing on the compute nodes cannot.

2.5 EDGE COMPUTING CAPABILITIES

The FLAME platform requires edge computing capabilities as a core requirement.

- There are two approaches that can be adopted: individual experimenters can provide edge devices or, alternatively for FLAME, the edge computer could be a shared resource provided by the infrastructure provider. Experimenters should be allowed access to their slice within OpenStack, where each customer VNF should live within the respective customers' slice.
- Additional switching will be required in the edge, next to the edge computing resource. Experimenters should have exclusive access to their slice, so this edge switch will need to support VLAN and SDN switching. It will need to be housed and sized appropriately, considering environmental and practical constraints as well as computing resource capabilities.

2.6 MANAGEMENT SYSTEM

A Management System should be developed to effectively manage and monitor the FLAME platform.

- Either bespoke or generic management systems can be deployed. Typically, some tools are used to monitor connectivity and link performance. Bristol has a tool that provides a dashboard showing connectivity and separate tools that show link performance over time.
- Major components of the system have their own dashboards for configuration and monitoring, for instance:
 - OpenStack Dashboard
 - Wi-Fi controller
 - SDN controller
 - LTE MME
 - LTE management system

- Optional – Configuration manager, e.g. Red Hat Satellite or Spacewalk
- Devices need to be sized and connected to provide adequate capabilities; for instance, some machines are housed in Kernel-based Virtual Machines (KVM) and others in OpenStack. Consideration should be made as to what facilities should be controlled by the platform and what is exposed to the customer. As a minimum, OpenStack and SDN control should be customer-facing but customers must only have access to their slice.

2.7 INFRASTRUCTURE SLICING

Each infrastructure slice is an OpenStack project. VLANs are uniquely assigned to projects aside from where multiple projects require connectivity to the same VLAN in such cases the VLANs can be setup as shared networks within OpenStack allowing multiple projects to access the same VLAN. Each OpenStack project is usually created with at least 2 VLAN networks by the administrator. The FLAME OpenStack project should have 3 VLANs a management network, an SDN control network and the data plane network. These networks are real city VLANs on the BRISTOLOPEN networking equipment and are trunked around the city, this allows non-OpenStack equipment to be able to be connected into a project. Examples are Clients' own kit plugged into city racks, or wireless devices connected to Wi-Fi AP SSIDs around the city which are configured to tag traffic with a project's VLAN id. Devices can get DHCP-configured according to the project VLAN network's DHCP settings, and will have connectivity to the project's Virtual Machines, subject to the project's internal firewall rules etc.

Project VLANS can route between each other, and to/from VxLan networks created in the project, by creating a virtual router and connecting the networks to it. Giving the router a gateway to an external network allows project networks to connect to shared BRISTOLOPEN resources, and shared resources in other projects, via the 'public' network if required.

If needed, project VLANs can also route to other subnets via the main BRISTOLOPEN firewall / router, by creating an interface on the firewall in the VLAN (all VLANs are trunked into the firewall). An SDN controller is deployed within the network that can configure VLANs and routing through the system and that will link to the overall FLAME orchestration functionality.

3 DEPLOYMENT AND COMMISSIONING

3.1 PLAN COVERAGE AREA

Deploying the FLAME platform begins by planning the connectivity areas within an environment, e.g. a city square, buildings and street furniture that require coverage. BiO recommends drawing these areas onto a map of the desired city locations and identifying potential street assets that can be utilised. The local government may be able to provide a mapping tool to identify AC connections, and street cabinet locations that are available.

Take notice of risks when selecting your coverage areas to minimise the chance of needing to revisit other potential coverage areas. Risks that should be considered include inaccessible locations, proximity to street level pedestrians, security risks to the hardware devices and physical conditions.

Third Party contractors can also provide surveys on the coverage within certain areas as recommendations for replicators.

3.2 ETHICS, LAWS AND REGULATIONS

Once the areas for deployment have been selected, the laws and regulations that apply to the specific selected areas must be considered. Consultation with the local government will be needed to ensure these sites are safe, secure and comply with public safety regulations.

If cameras are to be deployed, replicators will need to comply with data protection legislation to protect public data, and the local council will need to be alerted if they are to be installed on their land. Other considerations should be made to the unique ethical and regulation requirements which may vary in each city. Considerations should be made on the following factors;

- Does the local council/land owner have restrictions on the contractors that can be hired to deploy devices? For example; in Bristol, Bristol City Council has a list of pre-screened contractors that they allow to deploy devices within the areas owned by Bristol City Council.
- Hardware compliances.
- UMSUNG Codes for devices using electric supplies of local government.
- Public Health and Safety.
- Data Protection compliance.

3.3 RF PLAN / CONNECTIVITY PLAN

After deployment locations have been selected and confirmation from the relevant authority is received, the connectivity and radio frequency must be planned.

Decide upon the interconnects used across the network:

- Fibre – Fast and reliable interconnection. BRISTOLOPEN recommends fibre for the core of the network and ideally to the edge switches.
- Copper Unshielded Twisted Pair (UTP) – Can be used instead of Fibre if cabling budget is limited. However, the maximum theoretical performance when using copper UTP is significantly less than with fibre.

- Radio Link – If physical cables cannot be installed to connect a desired location to the network, a radio link may be used to provide a backhaul connectivity between the core network and coverage area. For example, one of the coverage areas could be across a river, making it costly to install fibre.

3.4 SWITCH LOCATIONS

Decide upon the locations for the core of your network. This could be for example: a secure building, cabinets or a rack within a SME's server room.

The edge switches should be deployed within cabinets, lampposts or secure buildings within the selected deployment areas. Replication cities will need to provide mitigating actions for risks posed by environmental and weather conditions. For instance, Bristol has issues with damp and cold issues when storing switches within cabinets. In other cities high temperatures and dust may be an issue, but this will need to be addressed by the replicators when assessing the risks of each location.

The edge switches should be installed as close to coverage areas designated for end user/end devices as possible to minimise latency for edge processing. For LTE deployments the length of cable from switch to LTE radio affects the coverage range of the radio.

3.5 LOGISTICS – CLOSING STREETS, AVAILABILITY, GENERAL PLANNING

Next, check the logistics of the deployment to reach out to site surveyors, land owners, etc. to confirm whether the logistics for the planned deployment are viable. Since each FLAME replicator's situation will be unique, the following considerations should be interpreted as guidelines.

Potential considerations:

- Street closures – Streets may need to be closed to install devices such as signal masts or Wireless Aps. Limitations to building work during certain periods of the year or bank holidays.
- Availability – When are engineers, equipment and contractors available?
- General planning – Plan the logistics in advance

4 RECOMMENDED SOFTWARE AND TOOLS

Infrastructure visibility and monitoring is key when deploying and supporting the FLAME infrastructure. Below is a list of recommendations for monitoring software.

Virtualization & Orchestration Technologies

- OpenStack [3] – it is a free and open-source software platform for cloud computing and virtual server lifecycle management.
- Apache CloudStack [12] – open source cloud computing software for creating, managing, and deploying cloud services.
- Docker [13] – it is an open platform for developers and sysadmins to build, ship, and run distributed applications, whether on laptops, DC VMs, or the cloud.
- OSM [4] – an open source Management and Orchestration (MANO) stack aligned with European Telecommunications Standards Institute (ETSI) NFV Information Models.
- ONAP [14] – an open source software platform that delivers capabilities for the design, creation, orchestration, monitoring, and life cycle management of VNFs; the carrier-scale SDNs that contain them; higher-level services that combine the above.
- ODL [9] – a modular open platform for customizing and automating networks of any size and scale.
- ONOS [10] – SDN network operating system designed for high availability, performance, scale-out.

Recommended monitoring software

- Nagios [15] – Provides connectivity monitoring. You can create a logical map of your network and setup ping connectivity tests to monitor the reachability of devices. Nagios can also be used to monitor the availability of the OpenStack APIs.
- Munin [16] – Provides server performance information including available storage, RAM, etc. Can be used to monitor your VM hosts.
- SmokePing [17] - To check network latency.
- OpenNMS [18] – Performance monitoring and event driven reporting.
- Ceilometer [19] or Jenkins [20] - To monitor OpenStack.
- Wireshark [21] - For packet analysis.
- Prometheus [22] - an open-source monitoring system with a dimensional data model, flexible query language, efficient time series database and modern alerting approach.
- Grafana [2] - Data visualization & Monitoring with support for Graphite, InfluxDB, Prometheus, Elasticsearch and many more databases.

Individual Key Performance Indicators (KPIs) are gathered by different management systems. Integration and reporting of these KPIs will have to be studied and considered depending on Network needs.

Management tools

BRISTOLOPEN recommends scripting the installation of infrastructure with a tool such as Red Hat Satellite [23] or Ansible [24] . This will increase the time of initial deployment but will allow to rebuild nodes with a couple of clicks in failure scenarios to minimise downtime. In addition, scripting the

installation will reduce the risk of human error due to repetition. Red Hat Satellite provides centralised monitoring and configuration management of infrastructure upgrades/updates.

5 INFRASTRUCTURE IMPLEMENTATION EXPERIENCES

This Section details the implementation of the FLAME architecture in the cities of Bristol and Barcelona. For both cities, details about the on-street deployment, the high level infrastructure topology, the FLIPS integration and initial experimental evaluations of the deployments are presented. The given information aims to facilitate parties interested in replication of the testbeds and serves as guidelines on how to implement the FLAME architecture.

5.1 FLAME INFRASTRUCTURE IN BRISTOL

5.1.1 Bristol Scenario and On-Street Deployment

Once replicators have access to the street level assets and have approval from the local government contractors or engineers, they will be able to deploy hardware to the street cabinets and lampposts.

Lampposts

Within Bristol, for the hardware devices to be installed on lampposts the approved contractor will need to install the devices using a cherry picker. For health and safety regulations, this cannot be done by BRISTOLOPEN engineers. Ducting from the street cabinet connects power supplies to the hardware devices on the lampposts.

Street Cabinets

Individual projects will utilise different parts of the infrastructure depending on their needs. Connectivity utilises a range of fibre, millimetre wave and microwave connections between radio sites and active nodes.

16 footway street cabinet locations dispersed around the BRISTOLOPEN fibre ring provide physical access to the 144-fibre cable.

These cabinets contain some or all the following;

BRISTOLOPEN Ancillary Devices

- Standard power sockets with in-cabinet circuit breaker
- Fibre termination unit
- Power over Ethernet
- Media Converter (SFP to RJ-45)
- Network management switch

BRISTOLOPEN Network Devices

- Wi-Fi Access points
- Bluwireless switch and/or relay

A typical street cabinet with the connectivity requirements is shown below (Figure 3):




Figure 3: Detail of the cabinet. BRISTOLOPEN testbed.

5.1.2 FLAME Infrastructure High Level Topology

The FLAME platform will be integrated with the infrastructure of BRISTOLOPEN in Bristol. Outlined in the diagram below (Figure 4) is a high level topology currently deployed within Bristol.

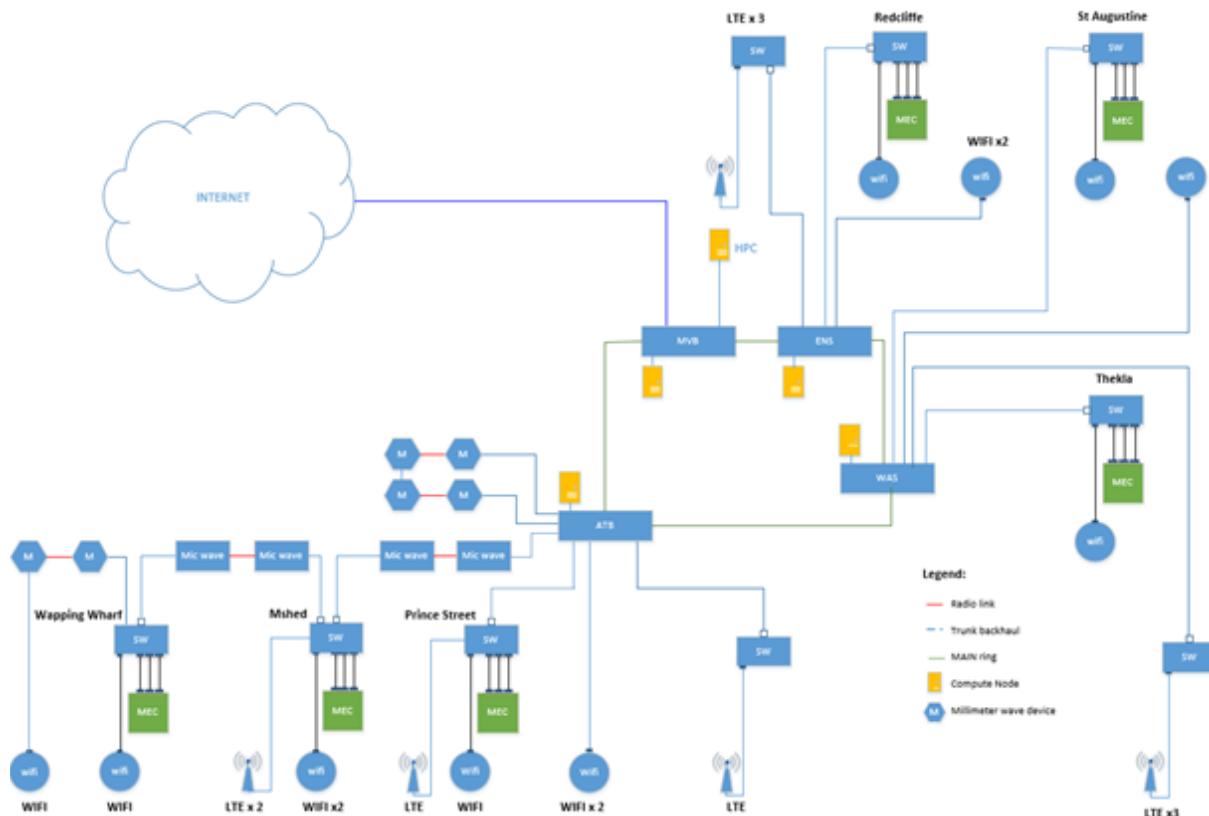


Figure 4: Bristol high level infrastructure

Compute Active Nodes

Bristol currently has 4 active compute nodes deployed in the city, which are connected by a fibre ring and provide connectivity between OpenStack and the wireless access points. They are in the server rooms of BRISTOLOPEN's key partners which are located in proximity to one another and the wireless access points across the deployment area. These nodes can be identified in the green 'main ring' in the diagram above.

Wi-Fi Access Points

There are 6 Wi-Fi access points that have edge computing currently. The edge computing capability will be integrated into OpenStack and become a sharable asset. There are an additional 6 WIFI access points deployed in the city that do not currently have any edge computing capability.

10 LTE base stations and Millimetre wave connections are also deployed around the city. These have no edge computing capability and are not currently included in the plans for FLAME.

These access points are installed in street level assets owned by Bristol City Council. Installations into these assets require BRISTOLOPEN complying with Bristol City Council contracting regulations and health and safety legislation i.e. working at dangerous heights. Additionally, equipment deployed into these cabinets require UMSUNG codes as they draw AC power from the city council.

5.1.3 Platform Integration

An experiment utilising FLAME concepts is deployed on BRISTOLOPEN's infrastructure. Figure 5 shows a simplified topology diagram. The two VMs (VNFs) shown on the diagram are located within an

OpenStack cloud which provides the compute capabilities needed to host the VNFs. Four Core switches can be seen at the centre of the diagram connected via a 40 Gbps fibre loop. The edge switches are connected to the closest core switch and provide connectivity to the network for the wireless access point and edge compute devices (APUs). The IPASO devices on the topology are 1 Gbps microwave links that provide backhaul connectivity to areas that cannot be reached with fibre.

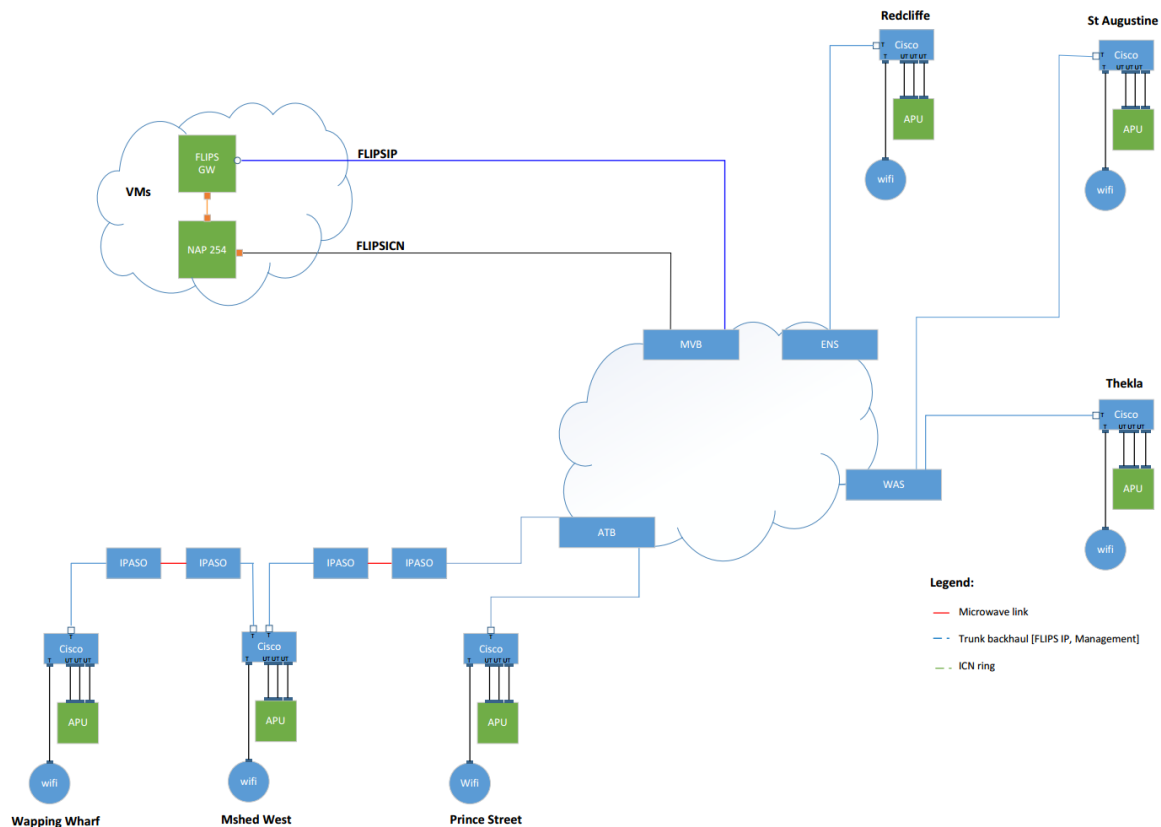


Figure 5: Platform integration at Bristol

OpenStack FLIPS Deployment

To setup an OpenStack instance as a router like the FLIPS GW VM in the diagram above, port security must be disabled on the internal and internet facing ports for the instance that should act as a gateway. Otherwise the *iptables* rules automatically added to ports to provide port security by OpenStack will drop traffic arriving on the internal interface destined for the internet. IPv4 forwarding must be enabled on the Gateway VMs Operating System and *iptables* rules must be configured to masquerade traffic exiting the internet facing interface. Additional *iptables* rules will need to be added to allow forwarding between internal and external interfaces.

5.1.4 Performance Measurements and Initial Deployment Results

Platform tests will be performed and validated against the infrastructure in stages. Initial tests will be done using Open stack and integrating with an APU unit deployed in an active node. Further functionality will be incorporated into the test plan as the project moves forward and results stored in the project folders. Wireless infrastructure is already deployed so will not be tested separately only as part of testing the full FLAME offering and initial scenarios.

5.2 FLAME INFRASTRUCTURE IN BARCELONA

5.2.1 Barcelona Scenario and On-Street Deployment

In Barcelona, the implementation of the FLAME architecture consists of (1) the on-street deployment that provides Radio Access Network (RAN) capabilities, (2) the Multi-Access Edge Computing (MEC) installations to provide light added value services close to the edge, and (3) the main DC deployment in i2CAT facilities. Main DC IT resources are used to provide heavy computational / storage services, e.g. high definition video content, video transcoding, quality of service and consumption analytics, as well as resource orchestration and management logic, e.g. OpenStack, ODL, DHCP servers, etc.

The on-street deployment consists of the wireless nodes mounted on lampposts that provide connectivity for user equipment over Wi-Fi. The Single Board Computers (SBCs) holding the wireless interfaces are integrated in adapted metal boxes containing also other crucial electronic and networking components, such as AC/DC converters (220V to low voltage for the SBC), fibre media converters (optics to electric), etc. The lampposts are connected via optical fibre with the FLAME edge infrastructure. In Barcelona, the edge infrastructure is deployed within a street cabinet, consisting of an edge server to enable ICN routing and providing VNF capabilities, as well as networking devices that aggregate traffic coming from the lampposts and also provide connectivity towards the main DC. The connection between the edge cabinet and the main DC has an intermediate hop in the IMI facilities at Glòries area, Barcelona. The connectivity has been guaranteed between the Glòries node and the cabinet setup at the Pere IV street via a fibre connection operated by IMI. Moreover, a private fibre optical link (owned by i2cat / IMI) has also been placed between the Glòries node and the main DC placed in the Omega building, Zona Universitària, Barcelona. In the following, each of the deployments (on-street, edge, and main DC) are explained in detail.

On-street deployment: wireless nodes on lampposts

For the on-street deployment, the Pere IV street in the district of Sant Martí was chosen. Within this street, a segment of around 400-500m hosts the deployment of the wireless nodes that provide RAN capabilities. Along this segment of the street, a total of 5 lampposts are manually picked in such a way that the nodes are more or less equidistant to each other and they follow a zig-zag pattern, switching from one side of the street to the other. Figure 6 shows an isometric view of the Pere IV with the FLAME street segment and the wireless nodes.

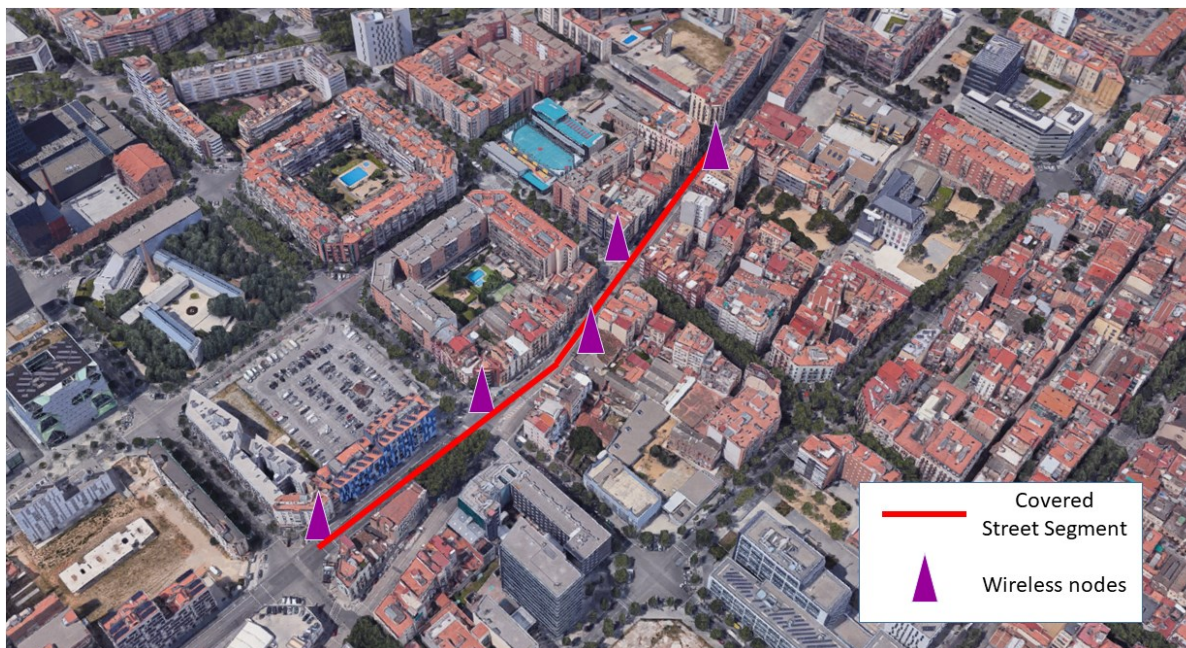


Figure 6: Barcelona FLAME on street deployment at Pere IV street

The lampposts provide both mains power and fibre connections for the wireless nodes. For the deployment of the wireless nodes, a third party was given the task of designing a casing that fulfils the following requirements:

- Weather-resistant
- Capable of switching from electrical network to optical network
- Capable of converting from 220V mains power (power line connectivity for standard "household" devices) to 48V to power the wireless nodes
- Providing a module that allows to reset the node remotely (hard reset)
- Providing a battery that activates if there is a loss of mains power, so the equipment can be turned off safely
- Providing fans for ventilation, to keep the temperature in the casing below any critical threshold

Based on these requirements, a casing and the internal components were designed. A look into the casing is provided in Figure 7), showing from left to right the remote reboot module, the media converter, the AC/DC converter, and the battery module. The empty space at the rightmost part of the casing is reserved for the Gateworks Ventana (GW) 5410 SBC [25], the platform that hosts the wireless interfaces.



Figure 7: View inside the box containing the wireless nodes deployed in Barcelona

The GW Ventana 5410 is equipped with either 2 or 3 wireless network interfaces of type WLE650V5 (Compex), implementing IEEE 802.11ac standard with backwards compatibility for the IEEE 802.11 a/g/n standards. One of these interfaces is always used for the RAN, i.e. it is used to instantiate wireless Access Points (APs), whereas the remaining 1 or 2 interfaces are used to provide wireless backhaul connectivity from each lamppost to its neighbours (thus nodes at the edges only require one backhaul interface, whereas intermediate nodes require two interfaces to point in two directions). For the RAN, omnidirectional dipole antennas are used, while for the backhaul directive panel antennas are used (both supporting 2x2 MIMO). The SBC also has two Ethernet ports that are both connected to the fibre media converter which enables a wired connection over fibre to the edge cabinet.

Edge deployment: cabinet server and networking devices

As stated above, Barcelona infrastructure includes an edge computing server and a switch, both placed in the cabinet at Pere IV street. The edge server offers application developers and content providers cloud-computing capabilities close to the end users. In principle, having services closer to the end user will improve the user experience [26]. Just as an example, resources on the edge computing server might be used for supporting the following: video analytic applications, location services, IoT, augmented reality applications, optimized local content distribution and data caching. In the context of FLAME, a portion of the edge server resources should be allocated for the instantiation of Network Attachment Points (NAPs) in the form of Virtual Machines (VMs) - one per lamppost. NAPs are mandatory elements to realize FLAME routing solution. In Barcelona deployment, the edge cabinet server is a 12 core CPU mini-tower server (shown below in Figure 8) with 128 GB RAM and ~2 TB of storage capacity. This machine has been registered as a compute node into the OpenStack controller hosted in the main DC. The Barcelona infrastructure setup represents a cost-effective infrastructure installation where a single cabinet server is assigned per several lampposts. Secure data and control / management communication lines between the main DC and the cabinet server are established as stated previously.



Figure 8: Cabinet server

Besides the cabinet server, the FLAME cabinet router (Cisco ASR920, Figure 9) is mounted in the street cabinet. It provides enough ports to connect from one side lampposts to the edge server and from the other side main DC into the FLAME street setup. Each fibre coming from a lamppost terminates into a Gigabit SFP at the Cisco router side. The FLAME setup, i.e. edge server and the Cisco ASR920 router with fibre connectivity to the wireless nodes on the lampposts, is connected to IMI over a 10 Gbps fibre.



Figure 9: Cisco ASR920 series

According to Cisco, the ASR 920 series is a full-featured converged access platform which provides a comprehensive and scalable set of layer-2 and layer-3 VPN services. It offers high throughput and low power consumption which makes it ideal for mobile backhaul, business services, video and data applications. The mentioned features correspond very well with the FLAME requirements.

The FLAME cabinet is connected to the main DC via a private network owned / operated by IMI. This network consists of two segments: 1- Optical network which connects Glòries node to the FLAME set up in the Pere IV street and 2- an optical network with maximum capacity of 8 lambdas (each supporting 10 Gbps) between the Glòries node and the main DC. The initial capacity considered on the optical segment is 20 Gbps but depends on the FLAME needs an adequate bandwidth will be allocated for the experimentations.

Main Data Centre

The FLAME street deployment in Barcelona is connected to the Data Centre, managed by i2CAT. The equipment hosted in the latter is used in a non-exclusive manner for the FLAME resources. It consists of computing devices, hosting a production-level virtualisation environment and the necessary networking devices that interconnect with the edge deployment through the IMI intermediate hop.

Specifically, the virtualisation environment consists of three servers running OpenStack Ocata, ranging from 32 GB to 96 GB RAM, from 4 to 6 cores and ~2 TB of storage capacity (Figure 10, Figure 11). The nodes exert the actions of controlling, computing and storing. The production-ready environment is defined following some of the best practices, such as high availability and redundant storage. To meet the former, the aforementioned servers provide different and/or replicated functionality. The control

plane of such cluster for virtualisation is architected in a way that punctual and/or localised faults can be overcome.





Figure 10: DC server (sample)

SuperServer SYS-5019S-MN4

(Top View – System)



© Super Micro Computer, Inc. Information in this document is subject to change without notice.

Figure 11: DC virtualisation cluster nodes

Moreover, the storage for the data required for the FLAME operation is distributed in a redundant way to overcome disk-related failures and maintain operation. The defined cluster manages all computational and storage resources over the main DC and the edge server.

Overall, the DC resources are utilized to instantiate and host functionality such as the control plane management for the wireless nodes (ODL), required storage capacity for media server content, potentially a number of DHCP servers to expose access to external users, and all the Network Access Points (NAPs) as required by the ICN routing schema. In Subsection 6.1.2.3, a high-level overview of these components is given.

More details about the OpenStack installation are provided below.

OPENSTACK: SERVER ARCHITECTURE

OpenStack components are mainly deployed in a Hyper-converged way. It means that some servers are working with controller, networking, storage and compute tasks. This architecture provides these benefits:

- A high available OpenStack installation can be deployed with only 3 physical servers.
- All hardware resources are fully used.

Figure 12 shows the major components of OpenStack and its deployment over the three main servers (IAS0 to IAS2) and the edge server (indicated as CAB1 in the figure). Since the CAB1 is located in the network edge, it will not act as a controller.

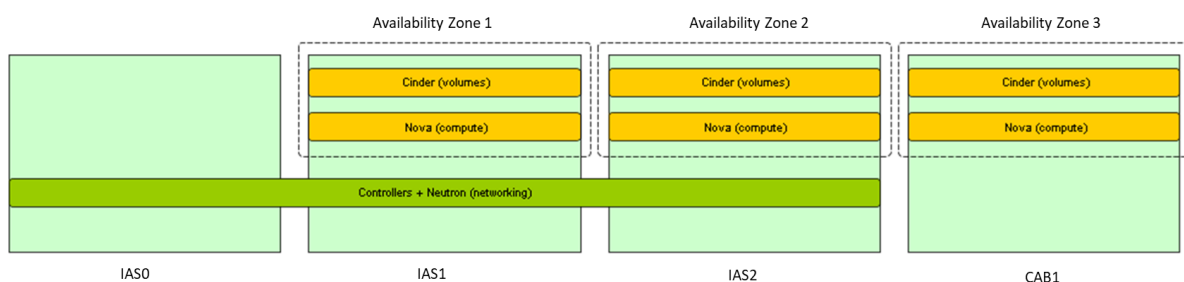


Figure 12: Deployment of the OpenStack components in the Barcelona main DC

In order to assure that all tasks are working with the proper resources, we use Linux containers for the controller/networking tasks, solid state disks for the storage and Linux KVM virtualization for the compute nodes. Availability zones are optional and can be useful in order to achieve better network performance. It gives for example the possibility to choose that virtual machines and their volumes are located into the same server node. We use OpenStack Ansible scripts to configure and deploy services over the physical servers. It should be noted that the specific configuration of the architecture is subject to changes in the future in order to best accommodate experimenter's services.

OPENSTACK: NETWORK ARCHITECTURE

We split the several networks using VLANs that operate with the different nodes in OpenStack. Figure 13 shows how:

- Different networks are separated into different VLAN (60 to 63)
- Different interfaces (1 Gbps Ethernet) are bounded providing:
 - High availability
 - Better network performance

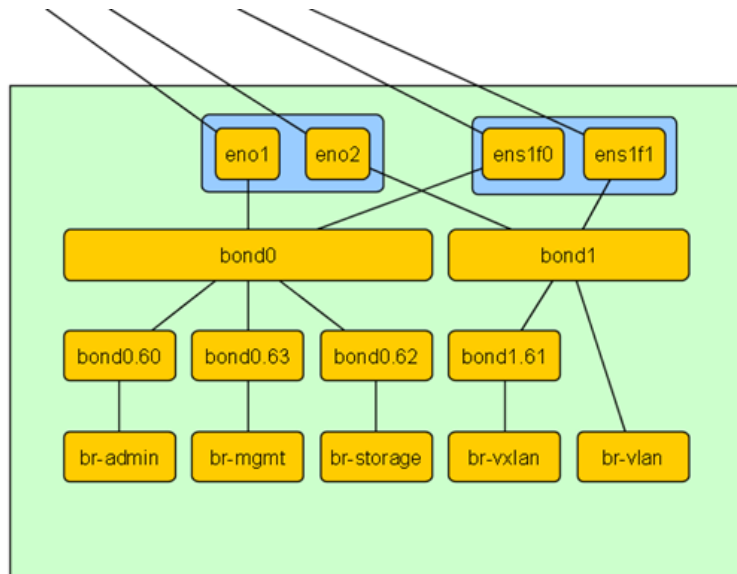


Figure 13: Interface and bridges bonding across the OpenStack servers

VLAN bridge connects directly to bond1 because it is where external VLAN will be dynamically defined. Figure 14 shows the whole network scheme. Note that main network switches are connected in a trunk mode, assuring the connectivity between the different OpenStack components.

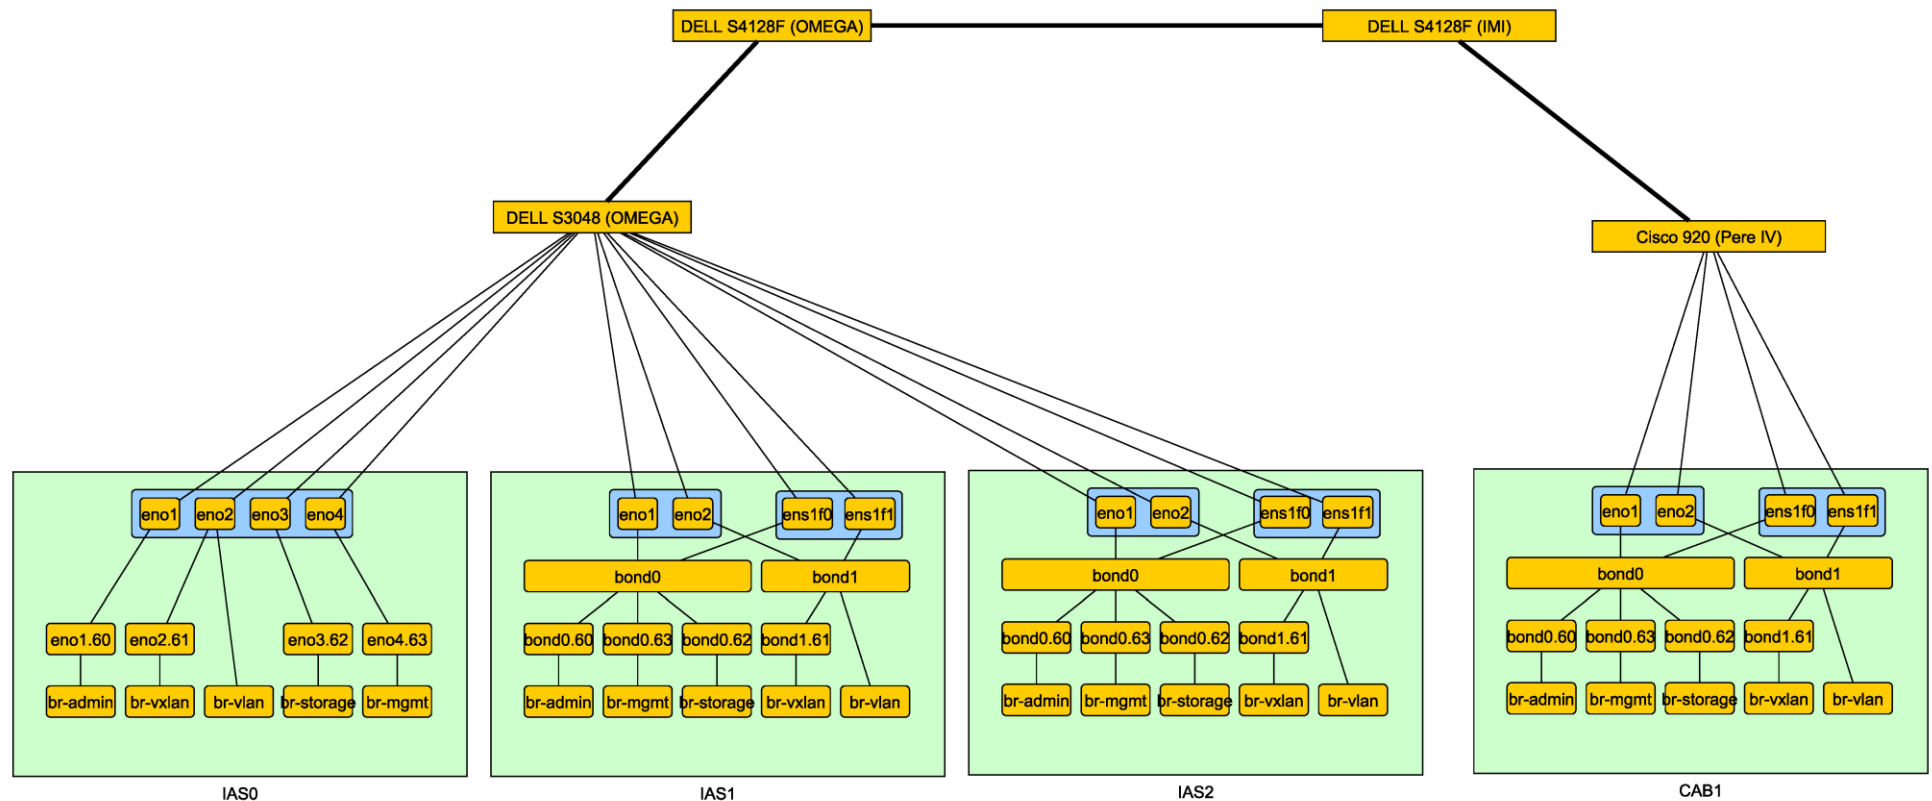


Figure 14: Connectivity between servers in the OpenStack cluster

5.2.2 FLAME Infrastructure High Level Topology

The Barcelona FLAME infrastructure consists of three main sites that are interconnected with a private network:

1. The Omega building close to the i2CAT premises hosting the main DC
2. The on-street deployment in Pere IV that hosts the edge equipment and the wireless nodes
3. The IMI premises that host networking equipment and serve as concentration point of the fibre connections from the Omega building and the street deployment.

Figure 15 shows a high-level topology overview depicting the three main sites.

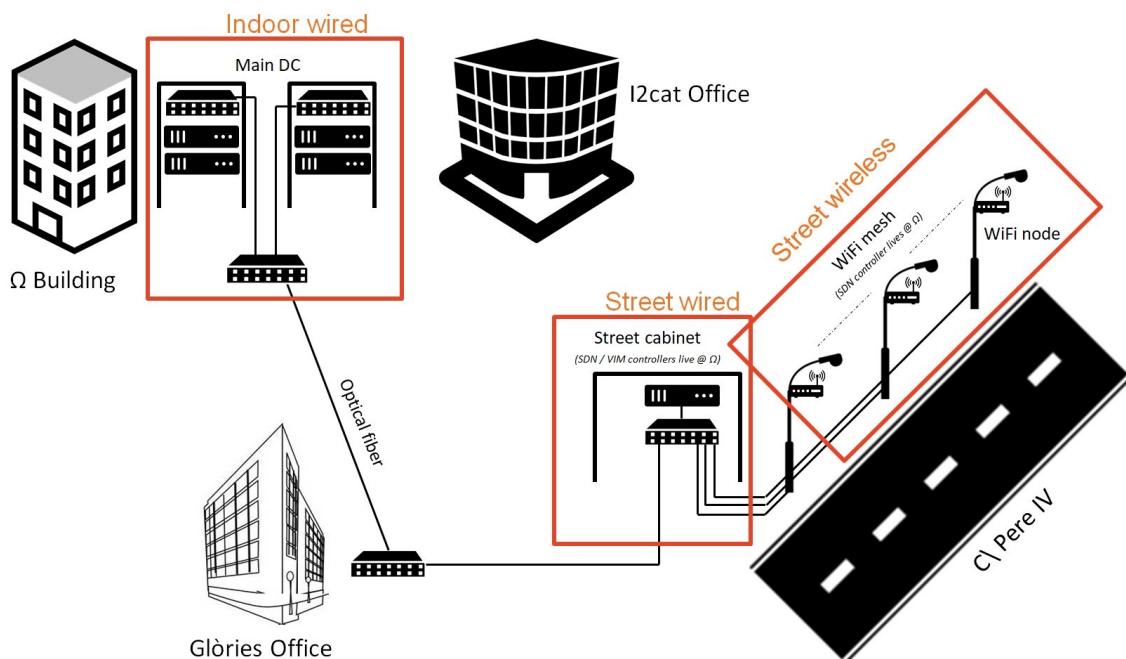


Figure 15: High-level view of the FLAME infrastructure in Barcelona

In the upper-left corner, the Omega building that hosts the main DC infrastructure can be seen. The main DC is composed of three servers connected to each other in the form of star topology (for the moment). The presented switch in the Figure 16 is a switch stack of two switches, which gives us a level of failure tolerance. However, to improve the main DC failover resilience, we plan to update the main DC topology in near future. The main DC is accessible over a fibre optic link from the i2cat office, but also supports external access via the Internet.




Figure 16: Network diagram of the main DC in Barcelona

In the lower left part of Figure 15, we see the IMI network that connects the Omega building with the on-street deployment. While the connection from Omega building to IMI happens over an optical network with maximum capacity of $8\lambda \times 10$ Gbps, the connection between IMI and the on-street deployment goes over a 10 Gbps wire to the Pere IV cabinet, where the FLAME cabinet setup is located. The on-street deployment is depicted on the right-hand side of Figure 15, also showing the fibre connections between the lampposts and the cabinet.

Figure 17 shows the different networking elements, capacities and the three sites of the Barcelona FLAME deployment discussed previously.

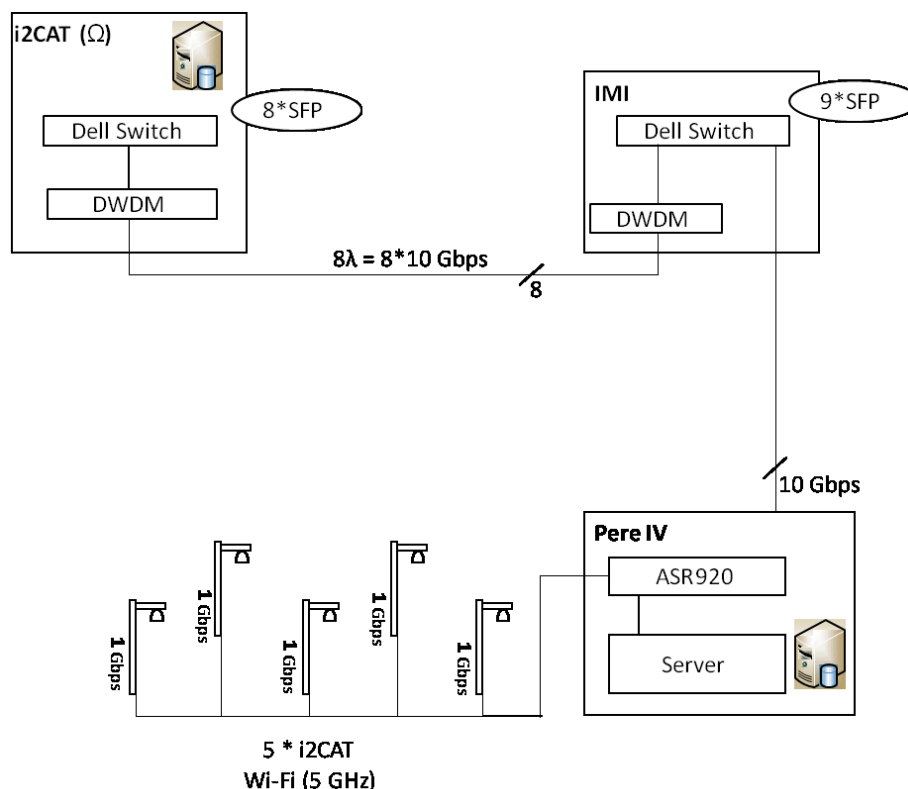


Figure 17: Network connections and networking equipment

5.2.3 Platform Integration

In order to enable FLAME to work on the Barcelona deployment, a planning is performed that maps the infrastructure requirements of FLIPS (as the main routing engine in the FLAME solution) to the actual city infrastructure. For FLIPS to work, there is a list of basic components that need to be provided:

- Computing resources in the main DC and/or edge to host services, such as OpenStack, NAPs, SDN controllers and/or media services
- Network connectivity between the main DC, edge and wireless nodes, which are capable of hosting several subnets to separate data, management and SDN planes
- SDN fabric to support the FLIPS routing

Barcelona deployment represents a cost-effective city installation where the FLAME solution could provide the significant leap forward for media delivery supporting personalized, interactive, mobile and localized (PIML) workflows. Leveraging on 5G-enabled programmable infrastructure, FLAME advantages such as faster access to media and services, lower latency and higher personalization of the experience through closer media processing will be offered through virtualized resources at the main and edge DC. This creates room for a significant reduction in the overall costs while ensuring fast availability of services towards end users. In particular, unlike the Bristol deployment where the hardware installation per lamppost is required to enable FLAME offerings (without providing any extra computing capacities for other added value service like content caching), resources on the general-purpose server mounted on the cabinet in Barcelona create a virtualized environment where NAPs as well as other added value services can be instantiated on demand. This will significantly reduce the installation cost (CAPEX) as well as maintenance and operational costs (OPEX). To ensure the duplicate

of Bristol deployment, in Barcelona setup one NAP (in the form of VM) per lamppost will be hosted in the edge cabinet. Of course, this imposes high computational resource demand on a single cabinet server. Therefore, in such cost-effective scenarios, it is essential to plan the installation accurately to ensure longer operational lifetime. For example, in our current setup with one 12-core server, 5 virtual CPUs are needed just to run NAP VMs for 5 lampposts in the street. It leaves us with limited number of cores / CPUs to host further services or to host more than one simultaneous experiments. Nevertheless, by adding an extra server at the cabinet, this problem will be easily solved. This clearly stresses the need for a precise resource planning on installations like Barcelona.

The deployment of NAPs or any other (media) service for FLAME service offering is performed during the bootstrapping phase via the OpenStack controller that is deployed in the main DC. Further, the SDN controller responsible for configuring the SDN fabric that is required by FLIPS and DHCP servers for mobile end devices are placed in the main DC. Figure 18 depicts all the critical components of the FLAME architecture as they are placed in the Barcelona deployment. Note that VLAN subnets are used to connect elements that belong together, such as the SDN control plane, the FLIPS data plane, etc.

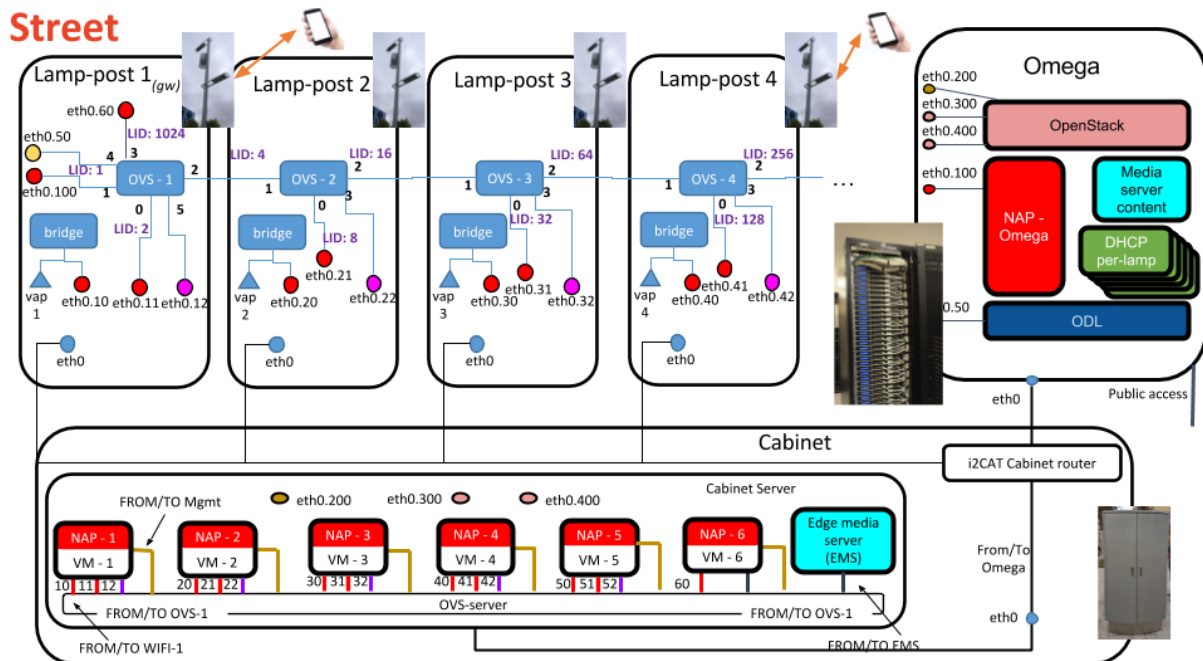


Figure 18: FLAME networking and component details for the Barcelona deployment

5.2.4 Performance Measurements and Initial Deployment Results

Before deploying the wireless nodes on the street, an evaluation is necessary to ensure that the chosen positions of the wireless nodes, i.e. the positions of the lampposts, can provide the desired on-street connectivity. Since the data traffic in the Barcelona deployment is carried over the wireless backhaul, it should be well performing and stable enough to assure connectivity between each lamppost and its direct neighbours. To evaluate whether this is the case in the Barcelona FLAME deployment, an on-street performance measurement has been carried out in which the link quality for different positions of the wireless nodes have been tested.

The experiments have been performed with two wireless nodes in order to determine the maximum distance between them that still permits a stable and high performing link. Given the natural signal

path loss over long distances, the data rates decrease with growing distance, as lower modulation coding schemes need to be applied, up to the point where the link becomes unstable. In the case of the Barcelona FLAME deployment, directional panel antennas operating in the 5 GHz band are used for the wireless backhaul communications. Directive antennas increase the potential range between two wireless nodes when comparing to the distances that can be reached with omnidirectional dipole nodes.

The experiments performed to determine the link performance at different distances consist in fixing the position of one node, setting up a wireless link with a second node, and performing a set of throughput experiments via Iperf [27] while choosing different positions for the second node.

Figure 19 depicts a bird's eye view of a Pere IV street segment. In the upper left corner of the overview, the first wireless node was set up at a height of approximately 2 meters, it did not move over the course of the experiments. The second wireless node was moved to 3 positions. Each link between the first and second node for these three positions is shown as a line in Figure 19. At each position a throughput test was performed. During such a test, the alignment of the sending node was varied along the horizontal and vertical axis in order to determine the performance in case the antennas are not directly aligned or even if there is no direct line of sight (LoS), taking into account that the directive antennas have a beam width of 25°.



Figure 19: Setup of the wireless nodes for the performance test in Pere IV.

The setup for the three experiments is shown in Figure 20. In the first experiment, at a distance of 40 meters between the two nodes, we observed a maximum throughput of around 220 Mbps (y-axis) over the course of the experiment (x-axis). The throughput decreased to around 150 Mbps when performing rotations of the directional antenna of the second wireless node. The rotations are performed up to 90°C within a plane in one direction, then turning back to 0°C and continuing up to -90°C. Each of these rotations is marked with the correspondent symbol underneath the x-axis at the time they were performed.

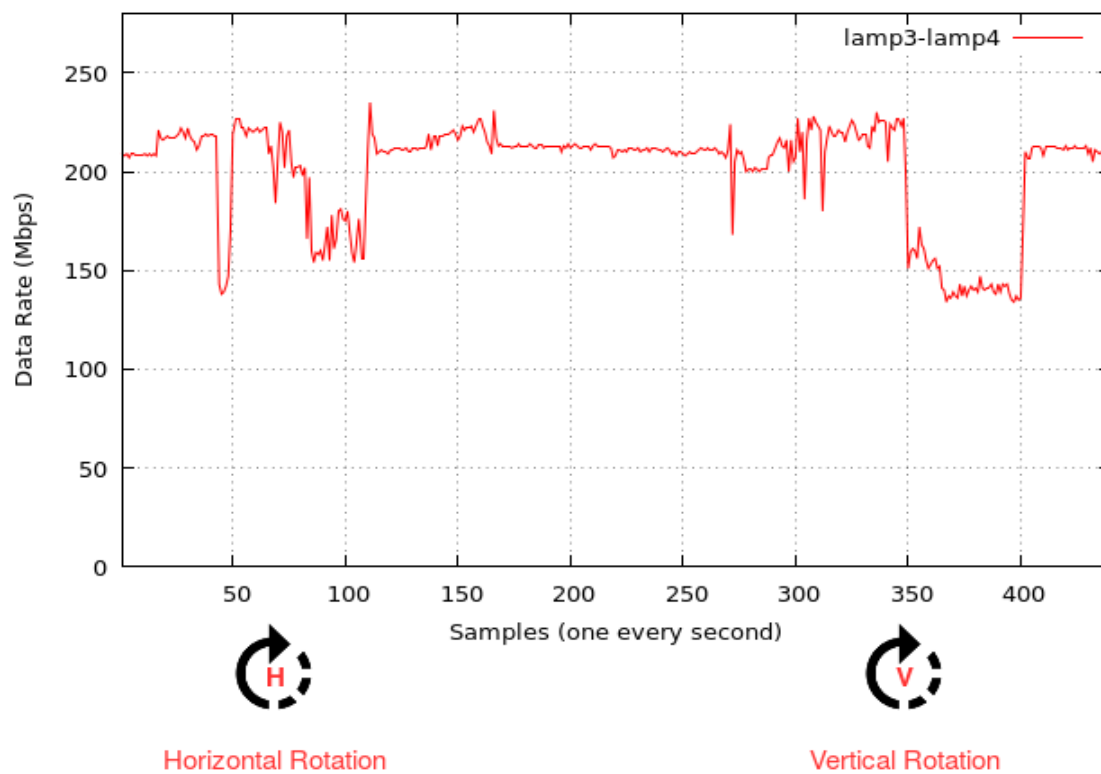


Figure 20: Results for the 40 m backhaul link

We observed that the link is very robust, as even aligning them at an angle of 90°C in the horizontal plane still allows for stable communications, even though at a lower transmission rate. We determine that this is possible, since we were working with a setup between a row of buildings that cause multipath propagation of the signal, effectively sending it to several directions. This assures that there is a decent signal at the receiver's antenna even when it is not aligned with the transmitter antenna.

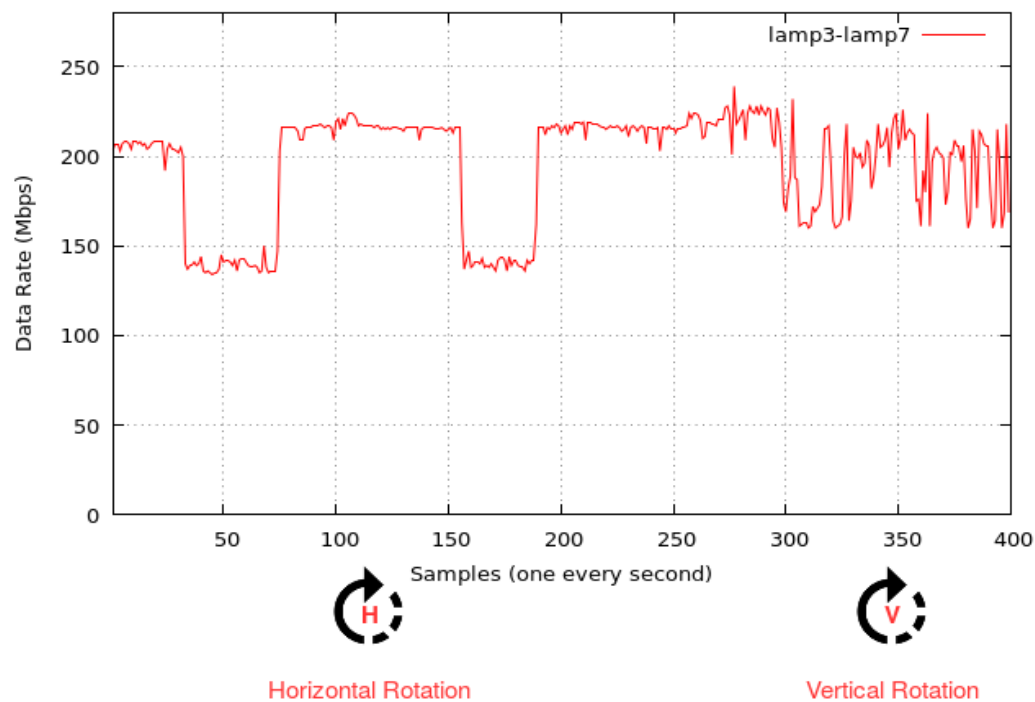


Figure 21: Results for the 107 meters backhaul link

We observed a similar behaviour at a distance of 107 meters (Figure 21): the maximum throughput achieved is close to 220 Mbps. As soon as we introduced the rotations to the receiver antenna, we observed a similar drop of the throughput as in the first position. During the vertical rotation, however, we measured a more unstable behaviour than before.

In the last experiment, we positioned the nodes at a distance of 205 m. Figure 22 shows the throughput results for this experiment, revealing that now the maximum throughput is around 200 Mbps and the link is more susceptible to variations in the alignment. The horizontal and vertical rotation cause a noticeable drop of the throughput, down to 50 Mbps. This is an important observation, as it shows how important the alignment of the two nodes becomes at larger distances.

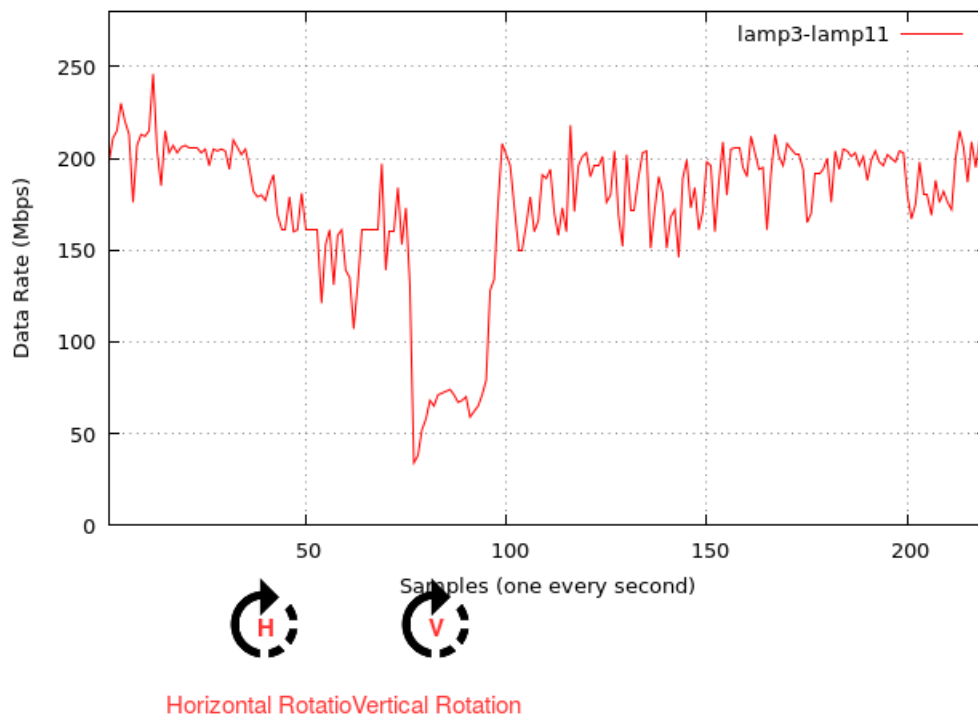


Figure 22: Results for the 205 m backhaul link

Based on the experimental results, we determined that the distance between two nodes with similar hardware should be at most around 100 m, a distance at which we can assume a stable and high performing wireless backhaul link if the two directive antennas are aligned correctly. Further, based on this decision, we determined the positions of the nodes for the FLAME deployment, choosing inter-node distances below the limit (whenever possible) and assuring that there is LoS. To assure latter, we choose the lampposts in such a way that they form a zig-zag pattern from one side of the street to the other.

Further aspects that need to be taken into account when doing an on-street deployment of wireless equipment involve potential obstructions of LoS between two transceivers due to foliage or vehicles. It is important to take into account that both foliage and vehicle as potential obstructive elements may vary depending on the time of day (e.g. large trucks used for delivery services are more likely to obstruct during the morning, when they are performing deliveries) or even the season of the year (e.g. foliage which in winter may not even exist, but may block the LoS during Spring, Summer and Autumn).

6 CONCLUSIONS AND RECOMMENDATION

In this document we have introduced the FLAME platform and discussed the replication process from both a business and technical perspective. The deployment and commissioning has been provided and the tools required to monitor the platform have been discussed. Finally, the initial experiences and details of the replication process that is being performed in Bristol and Barcelona is also provided as an example that can be followed by replicators.

7 REFERENCES

- [1] D3.5: FLAME Technology Roadmap v1. 20 November 2017
- [2] Grafana web page. Available at: <https://grafana.com/>
- [3] Openstack web page. <https://www.openstack.org/>
- [4] Open Source Mano. Available at: <https://osm.etsi.org/>
- [5] Real Decreto 1066/2001. Available at: <https://www.boe.es/buscar/doc.php?id=BOE-A-2001-18256> (in spanish)
- [6] Decreto 148/2001. Available at: <http://www.mldm.es/BA/PDF/DECRETO%20148-2001.pdf> (in Spanish)
- [7] Martin J. Reed, Mays F. Al-Naday, Nikolaos Thomos, Dirk Trossen, George Petropoulos and Spiros Spirou, "Stateless multicast switching in software defined networks", Online: <https://arxiv.org/abs/1511.06069>
- [8] Pica8, "PICA8: Programmable Internetworking & Communication Architecture, Infinite(8)", Online: <http://pica8.com>
- [9] OpenDaylight SDN controller. Available at: <https://www.opendaylight.org/>
- [10] Onos project. Available at: <https://onosproject.org/>
- [11] LoRaWAN specification 1.1. Available on online request: <https://www.lora-alliance.org/what-is-lora>
- [12] Apache CloudStack. Available at: <https://cloudstack.apache.org/>
- [13] Docker containerization platform. Available at: <https://www.docker.com/>
- [14] Open Network Automation Platform. Available at: <https://www.onap.org/>
- [15] Nagios tool. Available at: <https://www.nagios.com>
- [16] Munin tool. Available at: <http://munin-monitoring.org/>
- [17] SmokePing tool. Available at: <https://oss.oetiker.ch/smokeping/>
- [18] OpenNMS tool. Available at: <https://www.opennms.org/en>
- [19] Ceilometer. Available at: <https://docs.openstack.org/ceilometer/latest/>
- [20] Jenkins tool. Available at: <https://jenkins.io/>
- [21] Wireshark tool. Available at: <https://www.wireshark.org/>
- [22] Prometheus monitoring solution. Available at: <https://prometheus.io/>
- [23] Red Hat Satellite. Available at: <https://www.redhat.com/en/technologies/management/satellite>
- [24] Red Hat Ansible. Available at: <https://www.redhat.com/en/technologies/management/ansible>
- [25] Gateworks Ventana 5410 Datasheet. Available at: <http://www.gateworks.com/imx6-single-board-computers-gateworks-ventana-family/item/ventana-gw5410-network-processor>
- [26] Developing Software for Multi-Access Edge Computing, ETSI White paper No. 20, September 2017.
- [27] Iperf tool. Available at: <https://iperf.fr/>