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

D3.9: FMI Vision, Use Cases and Scenarios V2

01/03/2019

Steven Poulakos, Manuel Braunschweiler, Robert W. Sumner (DRZ);
Carlos Alberto Martin (Atos);
Marc Godon (VRT);
Gino Carrozzo, Francesca Moscatelli (NXW);
Julia Chatain, Fabio Zund (ETH);
Aloizio Pereira da Silva (UNIVBRIS);
Marisa Catalan, August Betzler (i2CAT);
Sebastian Robitzsch, Dirk Trossen (IDE);
Stephen C. Phillips, Michael Boniface (ITInnov)

This deliverable is produced within the middle of the FLAME project and provides a more informed description of the FLAME Future Media Internet (FMI) vision. The vision considers the evolving landscape of media services and technological paradigms that inspire 5G. The deliverable provides updated descriptions of FLAME partner validation scenarios and media service use cases, which leverage the content delivery optimizations provided by FLAME. The FMI vision described in this deliverable considers the three principle stakeholders: media service providers, platform providers, and infrastructure providers.

Work package	WP 3
Task	T3.2
Due date	31/10/2018
Submission date	01/03/2019
Deliverable lead	DRZ
Version	2.3
Authors	Steven Poulakos (DRZ), Manuel Braunschweiler (DRZ) , Carlos Alberto Martin (Atos), Marc Godon (VRT), Gino Carrozzo (NXW), Francesca Moscatelli (NXW), Julia Chatain (ETH), Fabio Zund (ETH), Aloizio Pereira da Silva (UNIVBRIS), Marisa Catalan (i2CAT), August Betzler (i2CAT), Sebastian Robitzsch (IDE), Dirk Trossen (IDE), Stephen C. Phillips (ITInnov), Michael Boniface (ITInnov)
Reviewers	Tomas Aliaga, Lamprini Kolovou (Martel) & Carlos Alberto Martin (ATOS)
Keywords	5G, Interactive Media, Future Media Internet, Experimentation, Trials

Document Revision History

Version	Date	Description of change	List of contributor(s)
V0.1	07/05/2018	Initial Outline	Steven Poulakos (DRZ), Julia Chatain (ETH)
V0.2	08/10/2018	Revised Outline	Manuel Braunschweiler (DRZ), Steven Poulakos (DRZ)
V0.3	12/11/2018	Merge of Contributions	Fabio Zund (ETH), Julia Chatain (ETH), Marc Godon (VRT), Gino Carrozzo (NXW), Francesca Moscatelli (NXW), Marisa Catalan (i2CAT), August Betzler (i2CAT), Carlos Alberto Martin (Atos)
V0.4	16/11/2018	Integration of parts; Removal of Platform Section	Steven Poulakos (DRZ)
V0.5	20/11/2018	Introduction and comments	Steven Poulakos (DRZ)
V0.6	26/11/2018	Structure changes and comments for improvements	Sebastian Robitzsch (IDE), Dirk Trossen (IDE), Steven Poulakos (DRZ)
V0.7	04/12/2018	Updated Sections	Carlos Alberto Martin (Atos), Marc Godon (VRT), Gino Carrozzo (NXW), Manuel Braunschweiler (DRZ), Steven Poulakos (DRZ)
V0.8	05/12/2018	Introduction and scenario updates	Marc Godon (VRT), Julia Chatain (ETH), Steven Poulakos (DRZ)
V0.9	06/12/2018	Updates to Infrastructure section	Marisa Catalan (i2CAT), Dirk Trossen (IDE)
V1.0	10/12/2018	Added abstract, executive summary; updates	Carlos Alberto Martin (Atos), Marc

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

		to FMI vision and VRT scenario; overall clean-up	Godon (VRT), Steven Poulakos (DRZ)
V1.1	11/12/2018	Conclusion and more clean-up	Steven Poulakos (DRZ)
V1.2	12/12/2018	Comments on FMI vision and VRT scenario	Carlos Alberto Martin (Atos), Marc Godon (VRT), Steven Poulakos (DRZ)
V1.3	12/12/2018	Updates to FMI vision and VRT scenario	Carlos Alberto Martin (Atos), Marc Godon (VRT), Steven Poulakos (DRZ)
V1.4	13/12/2018	VRT Scenario and review prep	Marc Godon (VRT), Steven Poulakos (DRZ)
V1.5	08/01/2019	Updates based on comments from Carlos and Michael	Carlos Alberto Martin (Atos), Michael Boniface (ITINNOV), Marc Godon (VRT), Gino Carrozzo (NXW), Manuel Braunschweiler (DRZ), Steven Poulakos (DRZ)
V1.6	21/01/2019	Updates to improve vision and homogenization	August Betzler (i2CAT), Marc Godon (VRT), Carlos Alberto Martin (Atos)
V1.7	23/01/2019	Updates and comments	Stephen Phillips (ITINNOV), Steven Poulakos (DRZ)
V1.8	24/01/2019	Updates in the infrastructure section	Aloizio P. Silva (UNIVBRIS)
V1.9	05/02/2019	Mostly clean-up of comments and formatting	Steven Poulakos (DRZ)
V2.0	15/02/2019	General review and comments	Tomas ALiaga, Lamprini Kolovou (Martel)
V2.1	27/02/2019	Handling reviewer comments	Gino Carrozzo (NXW), Steven Poulakos (DRZ)
V2.2	28/02/2019	Handling more reviewer comments	Marc Godon (VRT)
V2.3	01/03/2019	Final quality assurance	Michael Boniface (ITinnov)

DISCLAIMER

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

Project co-funded by the European Commission in the H2020 Programme		
Nature of the deliverable:	**R**	
Dissemination Level		
PU	Public, fully open, e.g. web	✓

CL	Classified, information as referred to in Commission Decision 2001/844/EC	
CO	Confidential to FLAME project and Commission Services	

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

EXECUTIVE SUMMARY

This deliverable is produced within the middle of the FLAME project, and it builds on "D3.1: FMI Vision, Use Cases and Scenarios v1" to provide a more informed description of the Future Media Internet (FMI) vision, use cases and scenarios. The FLAME FMI vision is presented, which envisages the landscape of media services in the next few years, considering the new technological paradigms and the principles that inspire 5G. The FMI vision is presented in the context of general user requirements and specific use cases that consider 5G capabilities. The vision also considers the three principal stakeholders: media services providers, platform providers, and infrastructure providers.

The "Participatory Media for Interactive Radio Communities" scenario utilizes the FLAME platform and infrastructure in Bristol. More specifically, the scenario uses the unique context of the Millennium Square at the centre of Bristol, not only from a technical perspective, but also from a media creation and broadcast perspective. Two radio communities collaborate: citizens and visitors to the square on one side and the public broadcast service provider on the other side. The square is a well-known gathering place and can sometimes be very crowded. In this context, the scenario tests the capability of deploying different strategies in using the resources at the edge for implementing an interactive and participative media experience. From a media creation perspective, the scenario tests new collaborative and contemporary media formats.

The "Personalized Media Mobility in Urban Environments" (PMM) scenario utilizes the FLAME platform and infrastructure in Barcelona. The PMM scenario aims at optimizing towards the edge the distribution of personalized Video on Demand (VoD) contents in order to allow the service provider to implement flexible and adaptable streaming services for its subscribers. The PMM service implementation makes a direct use of the capabilities offered by the FLAME platform, in particular for what concerns the automatic placement of functions and intelligent service routing.

The "Collaborative Interactive Transmedia Narratives" scenario utilizes the FLAME platform and infrastructure in Bristol. The scenario explores the use of local storage capabilities and on-edge computation to reduce latency and increase the resulting QoE. Assets related to the narrative are made available for download at the location where they will be consumed. There is also potential to utilize on-edge computation for low-latency state synchronization, making it possible to offer shared location-based experiences across multiple connected devices at a single location.

The "Augmented Reality Location Based Gaming" scenario utilizes the FLAME platform and infrastructure in Bristol. The scenario explores the use of local storage capabilities allowing the game players to download game assets associated with particular locations. The players may trade with a virtual economy at specific locations and grow gardens within larger coverage areas of the FLAME infrastructure. These capabilities make it possible to provide unique and personalized content for all the players.

This deliverable also describes the two city infrastructure locations on which the FLAME platform and validation scenarios are developed. The infrastructure locations are designed according to principles required to demonstrate key concepts of 5G architectures. These capabilities are then explored by FLAME. The capabilities offered by the infrastructures also provide important constraints on the validation scenarios. For example, the infrastructure utilized in Bristol has changed since the beginning of the project. The validation scenarios taking place in Bristol will now use the 5G-enabled architecture located at the Millennium Square in Bristol. This has motivated changes to the three validation scenarios that take place in Bristol.

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

TABLE OF CONTENTS

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

LIST OF FIGURES

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

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

LIST OF TABLES

ABBREVIATIONS

AP	Access Point
API	Application Programming Interface
CCTV	Closed-Circuit Television
CLMC	Cross Layer Management and Control
DC	Data Centre
DVB-T2	Digital Video Broadcasting – Second Generation Terrestrial
eMBMS	Evolved Multimedia Broadcast Multicast Services
E&M	Entertainment and Media
FMI	Future Media Internet
FMS	Foundation Media Service
HDR	High Dynamic Range
HFR	High Frame Rate
HPN	High Performance Network
ICN	Information Centric Network
IoT	Internet of Things
IP	Internet Protocol
LoS	Line-of-Sight
LTE	Long-Term Evolution (telecommunications standard)
MOS	Mean Opinion Score
MS	Millennium Square
NAP	Network Attachment Point
OTT	Over the Top
PIML	Personalization Interaction Mobility Localization
PMM	Personalized Media Mobility
PoP	Point of Presence
RAN	Radio Access Network
SDN	Software-defined Networking
SF	Service Function
SFC	Service Function Chain
TCP	Transmission Control Protocol
QoE	Quality of Experience
QoS	Quality of Service
UGC	User Generated Content
UHDTV	Ultra-High Definition TV
vCDN	Virtual Content Distribution Network
VSN	Virtualized Service Networks
VoD	Video on Demand
WCG	Wide Colour Gamut
WTC	We The Curious Museum

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

1 INTRODUCTION

This deliverable presents the second iteration of the developing vision of FLAME within the Future Media Internet (FMI). The FLAME project is motivated by the concept of optimising media content delivery through direct interaction with underlying network management functions. Today's common Over-The-Top content (OTT) distribution approach used for media services results in non-optimal resource allocations causing either overprovisioning costs for network operators or poor Quality of Experience for consumers. Boniface et al. [1] motivate the FLAME approach to content delivery that addresses emerging media demand trends through cross-layer integration and orchestration of media service function chains within an ETSI NFV compliant infrastructure. The aim of this deliverable is to describe the FLAME FMI vision and also to provide updated descriptions of FLAME scenarios and media service use cases, which leverage the content delivery optimizations provided by FLAME.

The **FMI Vision** provides the scope, which is used to select target vertical markets (e.g. TV, radio, gaming) for FLAME. As mentioned above, FLAME aims to optimize media content delivery by enabling deep interactions between media service providers and an underlying network which uses software defined networking and information centric networking techniques to provide features not seen in today's infrastructure. The **Scenarios** take the form of user stories, describing how users interact with media services deployed on the FLAME platform. The scenarios are characterized by demands required to enable specific experiences. Key characteristics in the experiences of the experiment participants considered by FLAME are Personalization, Interaction, Mobility and Localisation (**PIML**). Other demands, such as scalability or security, additionally contribute to the characterization of the scenarios. The **Use Cases** are media service and application use cases, which document functionalities required by the scenario, primarily with respect to the end user.

D3.1 FMI Vision, Use Cases and Scenarios v1 [2] provides a detailed description of the original scenarios and use cases. It provides a visual overview of the FLAME information model, which illustrates a scenario-centric relationship with the FMI vision, use cases, scenario evaluation and experiment trials. It also describes the FMI Vision aspects of the platform and enumerates the technical differences and specific capabilities of the FLAME platform compared to a conventional network and what benefits these differences bring. The resulting architecture is described in D3.3 FLAME Platform Architecture and Infrastructure Specification v1 [3].

This deliverable is produced within the middle of the FLAME project, and it builds on the initial Vision document to provide a more informed description of the FMI vision, use cases and scenarios. The outline is provided below.

1.1 OUTLINE

Section 2 "FMI Vision" provides the overall FLAME Future Media Internet vision. It describes the vision considering the new user needs and the respective technical requirements for the platform, which are satisfied by FLAME. The project has identified three primary key stakeholders: media service providers, platform providers and infrastructure providers (or owners). It also envisages the landscape of media services in the next few years, considering the new technological paradigms and the principles that inspire 5G.

Sections 3 - 6 describe each the following four initial scenarios selected by the FLAME consortium:

- ➔ Section 3 – Scenario: Participatory Media for Interactive Radio Communities

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

- Section 4 – Scenario: Personalized Media Mobility in Urban Environments
- Section 5– Scenario: Collaborative Interactive Transmedia Narratives
- Section 6 – Scenario: Augmented Reality Location-based Gaming

The above scenarios are represented by broadcast, gaming or transmedia vertical markets. Each scenario is introduced by describing its relation to the FMI vision. The scenario (or scenarios) and media service use cases are described, as well as details about the integration in the destination city. Updates and changes since the first version of the deliverable are highlighted. Each scenario is concluded by detailing its relation to the FMI vision.

Section 7 describes the city infrastructure on which the FLAME platform and validation scenarios are developed. The capabilities offered by the infrastructure provide important constraints on the validation scenarios. For example, the infrastructure utilized in Bristol has changed since the beginning of the project. The validation scenarios taking place in Bristol will now use the 5G-enabled architecture located at the Millennium Square in Bristol. This has motivated changes to three of the validation scenarios described above.

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

2 FMI VISION

This section envisages the landscape of media services in the next few years, considering the new technological paradigms and the principles that inspire 5G. This FMI (Future Media Internet) vision is focused on the user requirements and needs ([9] and [16]), which can be summarised as:

- **More information volume**, including the uplink and the downlink
- **Any device**, including fixed and mobile devices and any connected one
- **Anytime**, avoiding signs of congestion or delay even at the busy hour and at sustainable cost and enabling a continuous connectivity
- **Anywhere**, independently from the user's location or movement and including experiences linked to location. The consumer becomes mobile. Mobile devices are becoming the primary means for accessing E&M content. It is expected a continued growth of the media applications use in mobile scenarios.
- **Quality of service**, to prevent content frustration
- **Security**, for both end-users and content owners
- **Personalisation,** powered by data analytics and technology for better decision making

Most of the technical documents suggest a list of possible FMI use cases that can be deployed in a near future (and that consider 5G expected capabilities). The following ones are extracted from [9]:

- **On-site Live Event Experience**, which considers better experiences for the customers in large scale venues and crowded events, such as football matches. This use case includes large scale shared local experiences, which require computational resources on the edge.
- **User Generated Content & Machine Generated Content**. Users upload increasing amounts of self-generated content to the cloud and social networks. This trend is also verified by machines in the IoT paradigm.
- **Immersive and integrated media**. Other technical documents consider augmented reality and virtual reality [6]. These applications can take advantage of the low latency offered by 5G technology, which enables real time rendering of virtual objects, considering user movements.
- **Collaborative media production on the edge**, enabling the on-the-fly capture and sharing of video and audio (by means of cameras and microphones) and the work on these contents from multiple remote locations. Beyond the production, the availability of flexible resources on the edge also enables other capabilities for media producers, such as content ingest, local publication and local distribution.
- **Collaborative gaming**, enabling key improvements, such as more realistic experiences, collaboration among many simultaneous users and immersive environments. In this use case, the low latency is again a key characteristic.
- **Ultra-High Fidelity Media**. In this category, we can mention Ultra High Definition TV (UHDTV), which is not only more pixels (i.e., more resolution) but also better pixels. UHDTV can take

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

advantage of the larger throughput offered by 5G. UHDTV should be a mix of different audio-visual improvements: 4K, HDR, HFR, WCG and immersive audio.

These user needs and prototypical use cases imply technical requirements for the actors involved in the provision of media services and particularly the roles of platform and infrastructure providers. This includes aspects such as traffic volume (according to the Atos Look Out 2020+ report on Industry Trends for the Telecom sector, video and TV are expected to occupy 80% of the Internet traffic [6]; this trend is also striking if we pay attention to the estimation of video traffic in mobile networks: 72% of all mobile data traffic in 2019 [7]), latency and compute and storage resources distribution, which drives to new and flexible platform designs.

FLAME addresses accurately the characteristics of the FMI approach and it offers these features:

- Novel routing solutions to enable massive and efficient content distribution.
- High platform adaptability and configurability.
- Services linked to location.
- High interactivity and personalisation capabilities.
- Deployment of capabilities on the edge, which implies latency reduction.
- Instantaneous and multilayer monitoring. Metrics are crucial for media providers to keep users engaged. In the FLAME FMI vision this monitoring is used for the automatic orchestration of platform resources, considering the policies defined by the media providers.

The highly orchestrated and distributed platform enables the deployment of compute and storage resources on edge clusters. In this way, media service providers reduce the latency and optimise the local content delivery. This feature on the edge is key for some new services and scenarios, such as shared local media experiences, media content uploads (considering both professional and non-professional producers) and media content efficient distribution.

Inside the FMI ecosystem and its variety of actors, we see three principal stakeholders:

- **Media services providers.** Take advantage of these new paradigms to optimise the content distribution and the quality of experience. This approach goes beyond the current model of OTT contents based on the cloud: the vision explained in this document involves the infrastructure to ensure optimisation in each level.
- **Platform providers.** Provide a highly orchestrated and distributed platform, which implies a high configurability to fulfil service requirements in terms of compute, storage and connectivity capabilities.
- **Infrastructure providers and owners**. Infrastructure is not anymore just connectivity. It is connectivity and also multi-tiered data centres, where resources are flexibly orchestrated.

FLAME deliverable D2.1 [8] contains further information about the FMI ecosystem, the FMI actors and stakeholders and the role FLAME is playing.

It is expected that 5G integrates different network technologies, even in an adaptive way, depending on the number of users in a certain area, combining unicast, multicast and broadcast modes [9]. In this

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

sense, LTE already included a specification for the provision of broadcast services, named eMBMS (Evolved Multimedia Broadcast Multicast Services). This technology has been tested in similar conditions to the use case mentioned as "on-site live event experience". The existence of current broadcast networks (based on DVB-T/T2 standards in Europe) and the cost of the deployment of new cellular networks have driven to some hybrid proposals, as explained in [7] to reach a trade-off between the efficiency of broadcast technologies and the increasing demand of personalised content. These possible solutions show a new interplay field that integrates new and previous actors, including mobile network owners, infrastructure owners and platform providers.

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

3 SCENARIO: PARTICIPATORY MEDIA FOR INTERACTIVE RADIO COMMUNITIES

3.1 INTRODUCTION

Our focus on individual freedom has made some societal structures redundant. At the same time, a new search for identity and belonging has started. Now, people look for new communities wherein they get their status and confirmations. Digital borderlines are changing. They are not situated at the level of access to media equipment but are situated at the level of how to use them consciously and critically.

The Millennium Square is perceived as an important social meeting place, in the centre of Bristol, where stories about the future are initiated and shared. The term "Millennium" refers back the societal need for "The New Thinking for the New Millennium" [10] in the nineties. Three Millennium squares were built in the UK[1] in the light of that period.

The VRT scenario is created in the context of the search for new media formats where complex issues are addressed and both societal change and personal behaviour are positively stimulated. More specifically, the VRT interactive media scenario is developed in the context of a recently committed objective in helping to create more awareness on the UN Sustainable Development Goals and Climate Change challenges [11]. The discussions are taking place on the premises of Millennium Square (local cast); but we are also interested in how far away (geographically) these stories will pop up again (social and geographical reach of stories).

Interactive media allows users to explore and experience city stories and narratives. Users can pick up local posted media via the means of a physical and digital interaction. We use QR code technology to couple the material with the digital world. The media is retrieved from a digital inventory (e.g., archive). In this scenario, the inventory is hosted by the FLAME platform.

User-generated content or user commenting can be integrated. Contributors agreed, via consent, on the full understanding of the curation mechanism and ethics put in place.

We objectively verify QoE using different network QoS parameters using data acquisition, analysis and dashboard features offered by the FLAME platform. Two radio communities collaborate. Citizens and square visitors on one side and the public broadcast service provider on the other side. The square is a well-known gathering place and can sometimes be very crowded. In this context, the scenario tests the capability of deploying different strategies in using the resources at the edge.

In relation to the FMI vision, the media business is a competitive business fighting for eye-ball time, ear-time, user engagement, brand positioning, and to be the best in community building. Boundaries between broadcast networks and social network blur. Media become more agile, interactive, participative, and fast-paced. Feedback loops can drive highly agile and responsive creative media productions and experiences. Metrics become as crucial as the media itself. Highly responsive and high-quality delivery of valuable media can keep the user in flow and does not disturb the engagement.

[1] https://en.wikipedia.org/wiki/Millennium_Square

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

3.2 SCENARIO

The scenario is composed of two mechanisms deployed using the FLAME platform.

First, Figure 1 represents the process of associating QR tags and media at a location. Second, Figure 2 represents the process of accessing media at a specific physical location at the Millennium Square. During the deployment of the interactive media experience (Figure 1), media from an archive is linked to the QR-codes at the location. At the edge, the relationship is stored as meta data and the transcoding process is executed, and the different encoded versions are stored.

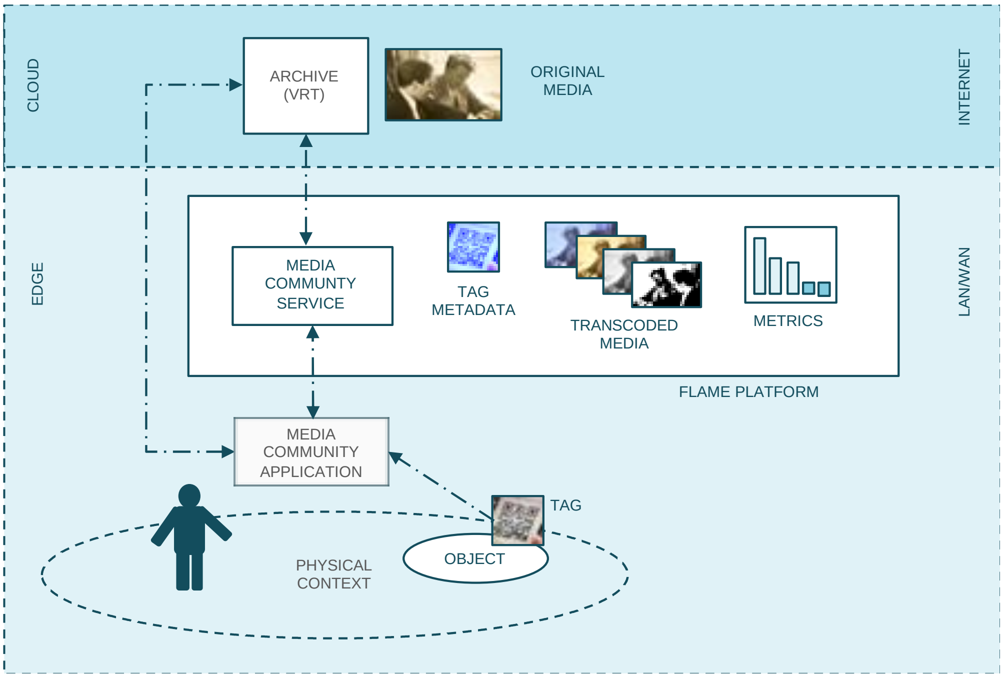

Figure 1 Associating QR tag and media at location (© VRT- Marc Godon, 2018)

When the user invokes the media service via reading the QR-code with his terminal (Figure 2), the media is personally delivered, with best effort to achieve a high quality of experience and user engagement. A quality strategy based on edge resource usage is deployed using the FLAME platform.

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

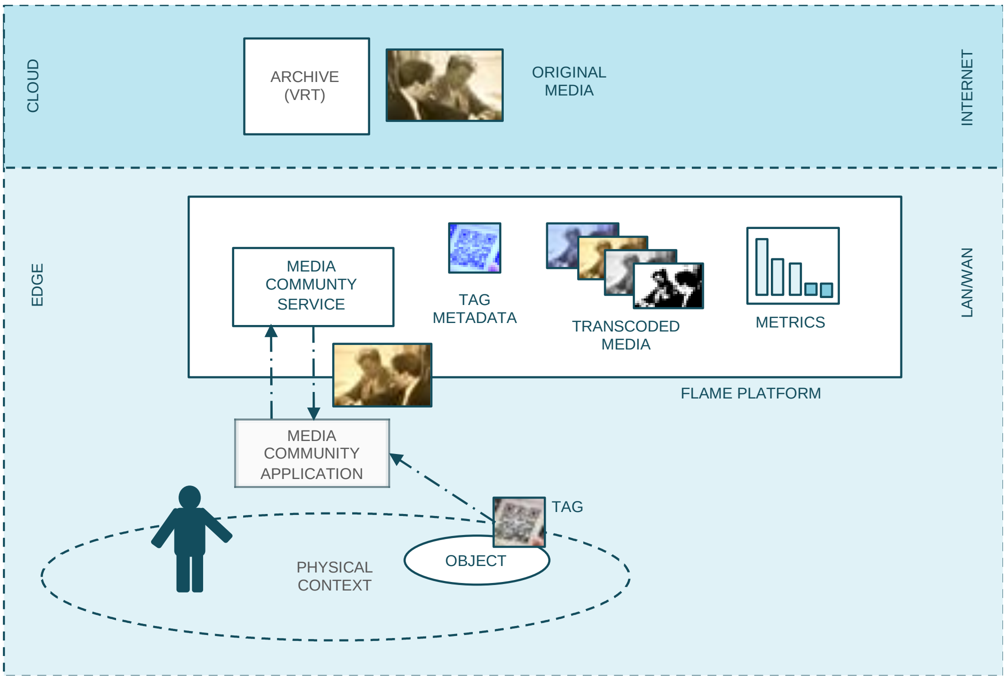

Figure 2 Access media at location (© VRT- Marc Godon, 2018)

3.2.1 Metrics

Metrics play a critical role in media. The quality of the media service delivery (related to the performance of the underlying network and the performance of the FLAME platform) can influence user engagement and overall appreciation [12] of the media experience (Flow [13]). Via the FLAME platform, we get access to these network and platform performance metrics. These performance metrics become more valuable if combined with the set of metrics commonly used in the media industry[2] (e.g. access to video objects and user viewing time). Using the same platform for defining and managing these metrics offers potential creative insights and input for new media design and contemporary media experiments.

3.2.2 Analytics and visualizations

Deployment of an interactive media experience in a physical space influences not only production planning and production management but will inspire in itself the subject and creation of the media content and the interaction design. A real-time heat map will plot different data dimensions (metrics). These visualizations can assist the media production personnel at the location in the management of the event.

[2] List of common metrics is developed with consult of VRT Study Center and are subject of research in the experimentation.

3.2.3 Participants and GDPR

Participants are not allowed to use the application outside the boundaries of the Millennium Square. All participants agree to the informed consent. The following types of data will be collected with the permission of the participants: demographic data, psychographic data [14], and sociographic data [15]. All participant-related data, which is stored in the FLAME platform, is anonymized[3] and stored conform the data management plan of the FLAME project, which conforms to H2020 guidelines and EU GDPR legislation. This data will be consulted only for dissemination purposes such as scientific publications in agreement with all FLAME partners. The validation test data model is subject to the bilateral GDPR agreements with UNIV BRISTOL.

3.3 MEDIA COMMUNITY SERVICE (MCS) USE CASES

In this paragraph, two use cases of the Media Community Service (MSC) are further elaborated. Ingest of new interactive media is described in the first use case (see Figure 3 and Figure 4). In the second use case, we discuss how the user request for access to media is elaborated by the collaboration of different service components deployed on the Flame platform on the edge (see Figure 5 and Figure 6).

Service components used in the Media Community Service use cases:

- MCS Controller (VRT): orchestrator or controller of the two use cases. This component is developed by VRT.
- Meta-Data Storage (ATOS): is the service component responsible for storing meta-data about the ingested media, the interaction tokens, and their relationships. This component is developed by ATOS and is part of the FMS suite of the Flame platform.
- Quality Analysis Check (ATOS). Service component, part of the FMS suite and developed by ATOS. Verifies media quality parameters of specified media objects.
- TRANSIX (ATOS). Service component, part of the FMS suite and developed by ATOS. Converts the encoding of a specific media object into a new set of media objects with specified encoding.
- Media Storage (ATOS). Service component, part of the FMS suite and developed by ATOS. Component specialized in and responsible for storing media.
- MPEG-DASH (ATOS). Component responsible for adaptive bitrate streaming over the Internet. Part of the FMS suite and developed by ATOS.

[3] https://www.ucl.ac.uk/legal-services/guidance/gdpr-anonymisation-pseudonymisation

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

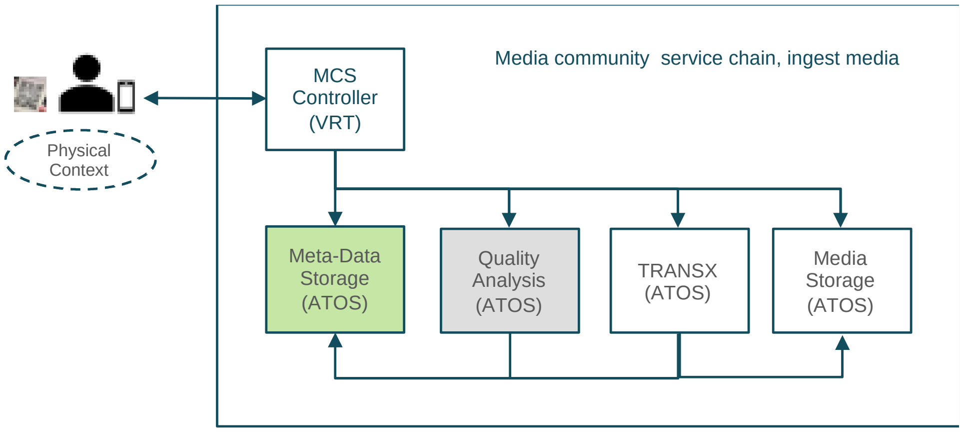

Figure 3 MCS chain for ingest of media

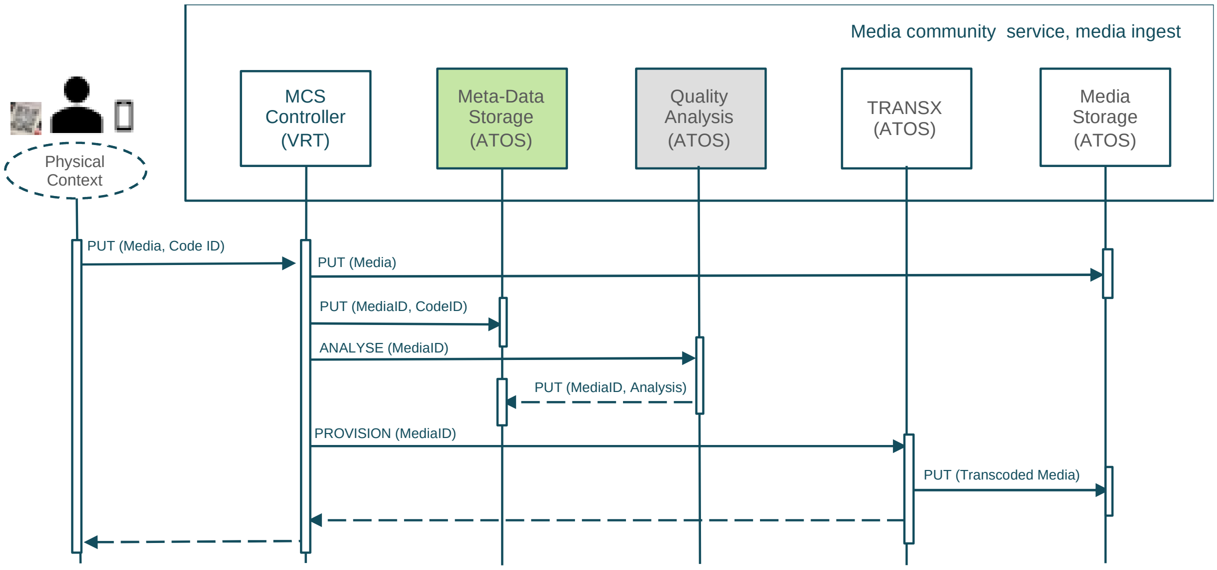

Figure 4 MSC scenario for ingest of interactive media (simplified)

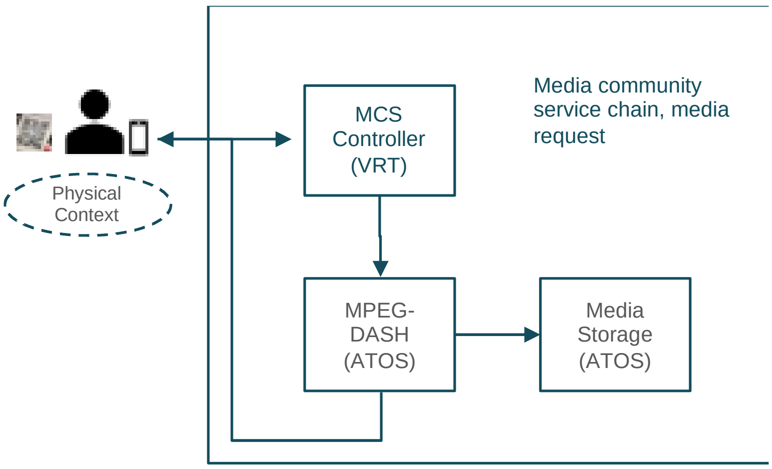

Figure 5 MCS chain, request media

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

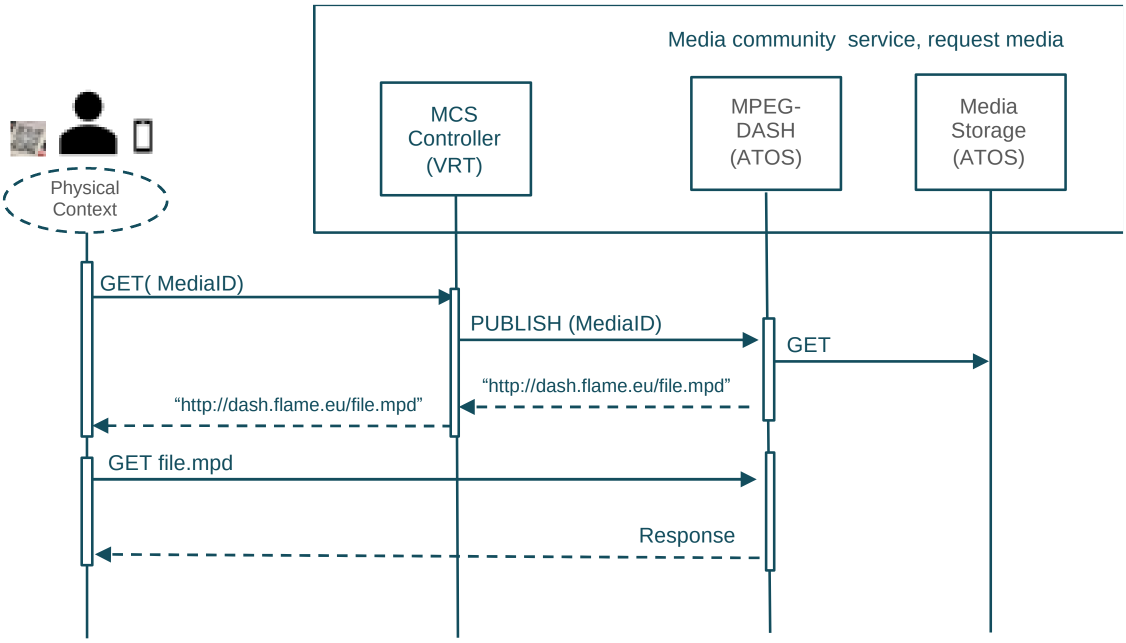

Figure 6 MCS request media sequence diagram (simplified)

3.4 INTEGRATION IN THE DESTINATION CITY

The landmarks on the Millennium Square are used as visual markers to guide users towards the touchpoints. QR-codes, which are the physical implementation of the touchpoints, are deployed nearby these landmarks. Via these touchpoints users actively interact with a locative distributed storyline about climate change and how citizens could join the debate or get inspired for action. Gamification techniques are used to disperse users over the square randomly. There is no fixed route. The QR-codes are more or less hidden so that people experience a kind of treasure hunt game. At the touchpoints, users can comment or share ideas.

3.4.1 The Millennium Square

Figure 7 gives a view on the north face of the Millennium Square. We can see (from left to right) the Planetarium, 'We the Curious' Science Centre, the BBC big screen, and stone benches.

Figure 7 The Millennium Square , at Bristol city centre, with the BBC big screen and Planetarium (© VRT- Marc Godon, 2018)

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

3.4.2 Landmarks

Table 1 gives an overview of the objects on the Millennium Square, which play the role of landmark in the scenario. This could mean that these objects not only have the role of a visual landmark, but they could also be related to the narrative of the mobile experience.

Table 1 Info box: Candidate landmarks at Millennium Square, Bristol.

 © VRT- Marc Godon, 2018	'Solar tree' to replace real tree in Bristol city centre (https://www.bbc.com/news/uk-england-bristol-30131406).
 © VRT- Marc Godon, 2018	Bill and Bob by Cathie Pilkington (http://aprb.co.uk/projects/all-projects/1999/bill-and-bob-by-cathie-pilkington).
 © VRT- Marc Godon, 2018	Aquarena by William Pye (https://www.williampye.com/works/aquarena).
 © VRT- Marc Godon, 2018	Kiosk at Millennium Square, Bristol.
 © VRT- Marc Godon, 2018	Planetarium (https://themathematicaltourist.wordpress.com/2013/08/03/millenium-square-in-bristol/).
 © VRT- Marc Godon, 2018	Benches in the middle of the Millennium Square, Bristol.

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

 © VRT- Marc Godon, 2018	Benches at the west-side of the Millennium Square, Bristol.
 © VRT- Marc Godon, 2018	Bollards to sit on at east side of Millennium Square, Bristol.

3.4.3 Physical Tokens

QR codes are printed as vinyl stickers glued on floor tiles or attached to poles or benches, as depicted in Figure 8.

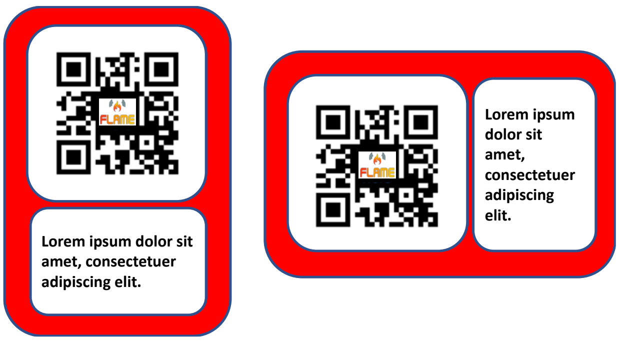

Co-funded by the Horizon 2020 Framework Programme of the European Union

Figure 8 Design QR tag, approximate size 10cm x 16cm © VRT- Marc Godon, 2018)

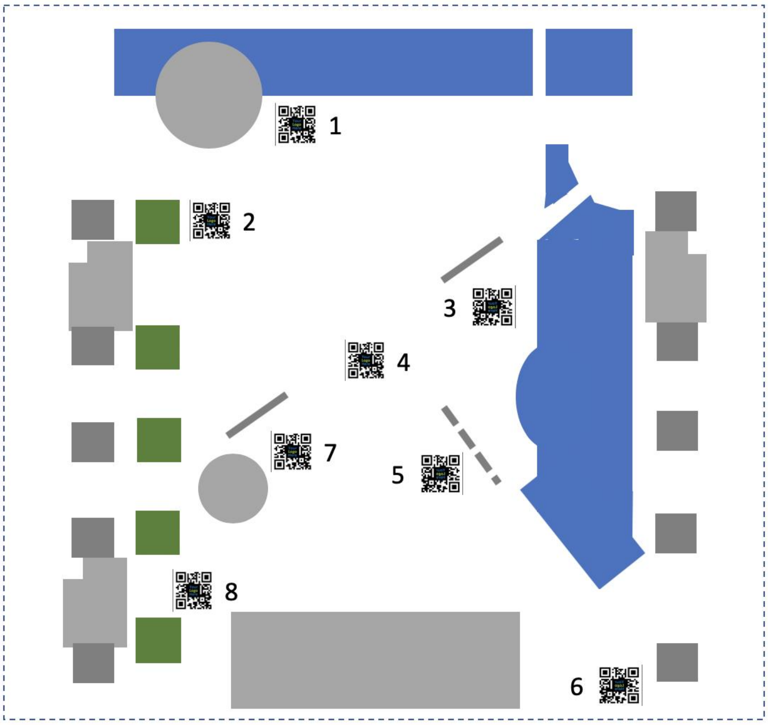

Figure 9 Outline of tags on the Millennium Square (dotted line = physical area of validation test, © VRT- Marc Godon, 2018)

Tags are located nearby the following landmarks (visualised in Table 1) as shown in top-down view of the Millennium Square (see Figure 9):

1. Planetarium
2. Solar Tree
3. Bench near Aquarena
4. Zenith[4]
5. Bollards
6. Bill and Bob
7. Kiosk or bench near kiosk
8. Bench

3.5 UPDATES AND CHANGES

The value proposition could be the result of the collaboration between partners from different business or different public service domains. The owner of the media rights can limit the rights of usage. Limitations can be expressed in relation to e.g. time, maximum number of views, or geography. Personalization is not only 1:1 but could also be defined regarding the characteristics of an audience or community. Interaction technologies such as location based human computer interactions (e.g. GPS

[4] Zenith by David Ward (http://aprb.co.uk/projects/all-projects/1999/zenith-by-david-ward)

based), QR codes, or NFC technologies should be part of a foundation interaction services framework.

3.6 RELATION TO FMI VISION

3.6.1 Converged Media 3.0

PWC outlines, in their recent publication "Global Entertainment and Media Outlook 2018-2022" [16], the drivers which establish a new Convergence 3.0. Ubiquitous connectivity: users are always connected and always on. Ever expanding the supply and diversity of content, experiences, and applications delivered directly and digitally to users. The mobile consumer: mobile devices are becoming the primary means for accessing Entertainment and Media content. Need for new sources of revenue growth. Value shift to the platform: platforms rather than publishers. Personalization: data analytics and technology for better decision making. In that same publication, the question: "What can a small Entertainment and Media player do to stay in the game?" is reformulated to: "What are the relevant capabilities to build relevance at scale?". The following capabilities are listed as being relevant: be a powerhouse of quality and engagement; aim at high-value and hard-to-reach audiences and deliver content and advertisement consistent with brand and user expectations.

3.6.2 Converged Media and FMI vision

Development of the scenario experiments has helped us to identify, in the role of public broadcast service provider, the following services supported by the FMI vision:

- ➔ Support for intelligent ingest of media in the edge. For example, generating metadata such as object detection, face detection, or in the context of the first level of quality control on the technical properties of the ingested media. This metadata can play a valuable role in the creative process. Semantic search will stimulate the creative process of editing and montage. It will allow for faster editing and thus more experiments. It will also allow for the detection of inappropriate content or give support for GDPR. This is especially valuable for mixing professional media with user-generated content, for example, by assisting the expert curation during a live event.

- ➔ Support collaborative production in the edge. Post, pre, and the central part of production could be computer aided in the network edge, which may support tasks such as creating scripts, capturing and storing media, and post production practices such as colour grading. Managing and editing different stream inputs before sending the resulting stream to a central office for further broadcast.

- ➔ Support for a local publication. It is possible to defer content personalization or adoption until the edge for reducing network communication cost. The publication may also target a local community using geographically defined networks.

  Support for local distribution. A local radio show provides access to local news. It would now be possible to provide access to media via new physical interactions. For example, visual object detection, QR-codes, NFC, community-created codes could be used to localize content. AR/VR interactions could also be provided.

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

4 SCENARIO: PERSONALIZED MEDIA MOBILITY IN URBAN ENVIRONMENTS

4.1 INTRODUCTION

This chapter reports achievements and modifications/extensions to the Personalized Media Mobility (PMM) scenario with respect to what was initially described in D3.1 [2].

The main goal of this scenario remains unchanged: i.e., to investigate how consumers participate and access personal media on the move. While we progressed on the execution of the PMM use cases, we decided to focus more on how service providers can serve personal media contents while the users are moving in the Smart City, e.g. to stream personal Video on Demand (VoD) and music on the go. Therefore, the PMM experimentation exploits the FLAME platform by setting up a personal media streaming service and intends to validate/evaluate this type of service on top of FLAME in terms of technical feasibility and business profiling, taking into consideration the involved stakeholders.

Two different narrative flows have been identified so far for the PMM scenario: "My Screen follows me" and "The Congress Eye" (i.e. urban scale personal media for public events). With respect to the scenarios initially proposed in D3.1, no major changes to the narratives and the related identified use cases apply and are reported in this deliverable. Though the PMM vision remains confirmed, the experimentation needs to consider the infrastructure capabilities of the actual target experimentation field (i.e. the FLAME infrastructure in Barcelona). Therefore, the current plan for experimentation of the "My Screen follows me" scenario does not include injections of live media contents from CCTV cameras into in the streaming server as initially envisaged and only VoD streaming from the users' personal libraries is being evaluated and validated. Moreover, at the time of writing this deliverable, the implementation of the "The Congress Eye" scenario results not feasible due to the location of the experimental infrastructure with respect to the major congress areas in Barcelona.

With respect to the FLAME FMI vision, the PMM scenario aims at optimizing towards the edge the distribution of personalized VoD contents in order to allow the service provider to implement flexible and adaptable streaming services for subscribers. The PMM service implementation makes a direct use of the capabilities offered by the FLAME platform, in particular for what concerns the automatic placement of functions and intelligent service routing. In particular, by experimenting the "My Screen follows me" scenario we aim at realizing a "follow-me" streaming service, exploiting the FLAME platform to adapt the service chain while the user is on the go, then placing/activating a streaming server component close to the user (i.e. in street cabinet) in order to reduce latency and offer better QoE.

4.2 SCENARIO

The validation and evaluation of the PMM scenario takes place in the city of Barcelona, where the FLAME platform has been deployed (see Section 7.2 for more details about the Barcelona infrastructure). The infrastructure comprises the core Data Centre (DC) in UPC campus area and the infrastructure deployed along Calle Pere IV acting an edge Point of Presence (PoP), whose access by the end users is foreseen via RAN nodes located on the street lamps.

Through the FLAME platform in Barcelona, we intend to investigate how FLAME can enable the distribution of personalized media contents in a Smart City environment while the user is on the go. In particular, our scope is to validate and evaluate the FLAME platform with respect to the PMM service orchestration and optimization, in order to investigate how service providers can serve the requested

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

media content in proximity to the end-user. In this direction, starting from the PMM scenario narratives and requirements described in D3.1, the PMM scenario experiment has been designed for being setup in different phases, involving an increasing number of end-users. The three identified experimentation phases and related scenarios are described below.

Scenario 1/Phase 1: PMM distribution in walking areas in Barcelona

This first experimentation phase is related to the scenario “My screen follows me” and involves a small group of end-users (up to 5 persons). In this “scenario 1” we aim to demonstrate the swipe of a media content from a fixed video/audio device at home (e.g. a TV screen) to a mobile device (e.g. smartphones or pads), adapting the service in order to guarantee the media streaming while moving in the Smart City.

In this scenario the PMM Follow me Media application must operate on top of the FLAME platform in order to instantiate the service via its North-Bound APIs exposed by the FLAME orchestrator and the CLMC. The FLAME platform is responsible for the service chain adaptation on the basis of specified service policies.

Scenario 2/Phase 2: PMM on aggregation areas of the Smart City

The “scenario 2” is the second phase of the “My screen follows me” experimentation and involves a larger group of end-users (from 10 to 50). As in the first case, in this trial the user swipes the media content from a fixed video/audio device at home to personal mobile devices and moves within the FLAME-covered urban area. If the user is in an aggregation area (e.g. shop, cafeteria, mall) the service chain is adapted, re-allocating the content caches to serve him/her.

Also in this case, the PMM Follow me Media application must instantiate the service making use of the FLAME platform APIs and the service chain adaptation/optimization is performed specifying proper lifecycle policies.

Scenario 3/Phase 3: PMM in digital signage posts

The third phase of the PMM experimentation deals with “The Congress Eye” scenario and involves more than 50 end-users. During this trial, users can consume video streaming (live and recorded) from various sources (e.g. cameras), placed at the Congress centre, on their personal devices while moving in the Smart City. The same library is available at digital signage posts installed at public places in the Smart City, allowing the swipe of the proposed media contents from them to the personal devices. Moreover, the digital signage posts can be used for delivering to citizens and delegates site-based media contents, e.g. video streaming from large public events in the Smart City.

Like in the previous two phases, the PMM Follow me Media application will be responsible for the service instantiation along with the policies submission, while the FLAME platform will orchestrate the service and performs the needed service chain adaptation/optimization.

As described above, while the scenario remains valid in terms of vision, its actual experimentation is not possible due to the specific logistic and coverage constraints of the FLAME infrastructure in Barcelona.

4.3 MEDIA SERVICE USE CASES

For realizing the above described scenarios, the PMM experimentation will involve FMSs provided by

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

ATOS and media service components provided by Nextworks combined in different service chains:

- **Origin Streaming Server:** this component serves the end-user's requests for streaming media content (VoD/Music).
  - **Opt#1 based on Symphony media server by Nextworks** with an embedded metadata DB and transcoding functions for pre-transcoding contents or performing adaptive streaming. It could be populated with recorded contents from live sources like the CCTV cameras.
  - **Opt#2 flame-fms-streaming**, works in cooperation with above listed FMSs.
- **Media Storage (flame-fms-storage)**, this component contains all the on-demand video and music contents and can be populated also with video recordings from the live sources (i.e. CCTV cameras).
- **Transcoder (flame-fms-transcoding)**, this component is used to transcode and/or condition (if needed) the original media content and to inject the transcoded content into the Media Storage.
- **Metadata DB (flame-fms-metadata)**, this component is used to store metadata information about media content stored in Media Storage.
- **Quality Check (flame-fms-quality)**, this component is used to retrieve information about stored contents' characteristics in terms of codecs, frame-rate etc.

Starting from these building blocks, the PMM scenario is realized through the deployment of two different service chains.

➔ **Service Chain #1**: this first service chain foresees the usage of just the Origin Server provided by Nextworks, which is an all-in-one media service capable of organising contents stored in a structured file system attached to the Server. The Origin Server will be responsible also for the pre-transcoding or live transcoding of media contents. Figure 10 depicts the interaction of the different building blocks. Figure 11 reports the sequence diagram for this service chain starting from the user's request until the delivery of the media content at the edge, in particular two different service flows are represented:

  - <u>Opt#1</u>: after the user's request to the PMM "Follow me" Media app, the user is redirected to the streaming server in the service chain capable of offering the best QoS/QoE, i.e. the one in the closest distance from the user. Then, the streaming server gets from its library the full quality version of the requested media content or a pre-transcoded version of it according to the network capacity. The user starts streaming the requested content.
  - <u>Opt#2</u>: in this second option, a pre-transcoded version of the requested media content is not available and the network capacity is not enough for streaming at a full quality. The streaming server, after getting the content from the media library, starts transcoding live the media content and serves it to the user.

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

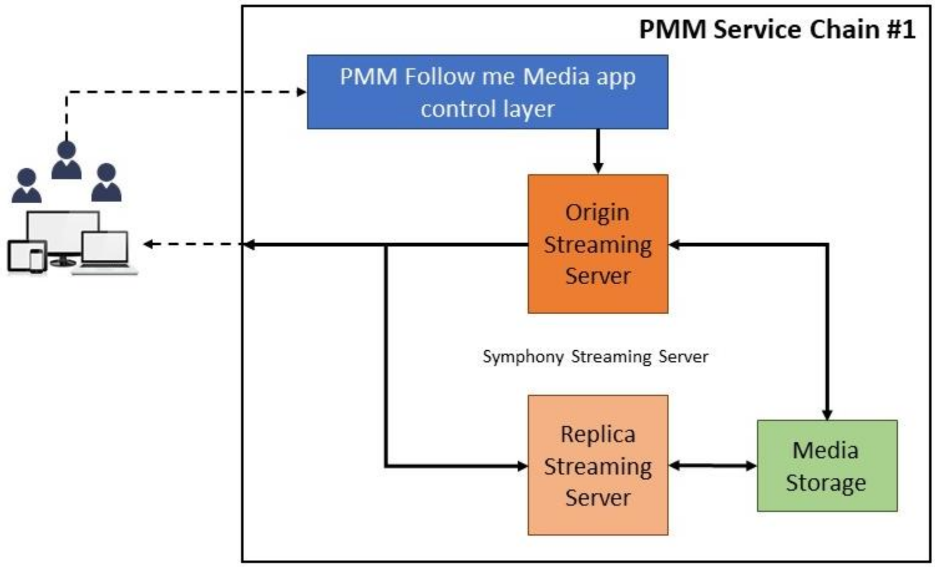

Figure 10: PMM service chain #1 media functions interaction

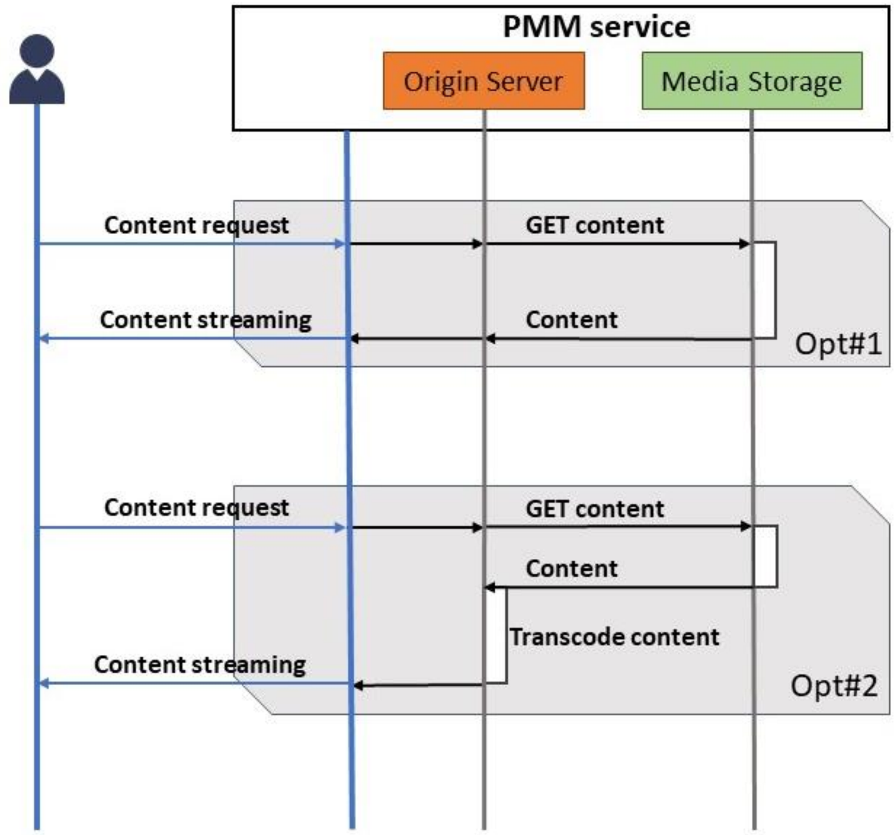

Figure 11: PMM service chain #2 simplified sequence diagram

➔ **Service Chain #2**: this second service chain includes the usage of the flame-fms-streaming server along with the correlated FMSs. Figure 12 depicts the interaction between such FMSs in this service chain, while the sequence diagram in Figure 13 reports the service chain workflow

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

starting from the users' requests to the media content delivery at the edge. As in the previous case, in Figure 11 reported two different service workflows:

- Opt#1: After the user's request to the PMM follow me Media app, the media content available formats in terms of quality are checked through the flame-fms-quality function. In this first case the most suitable quality for meeting the user's request is already available in the Media Storage (flame-fms-metadata + flame-fms-storage), the media content is injected in streaming server (flame-fms-streaming) and served to the user.
- Opt#2: In this second case, the most suitable format for the media content is not available, then the PMM "Follow me" Media app control layer is responsible for performing the transcoding/conditioning of the media content through the flame-fms-transcoding function. Once the content is properly transcoded, it is stored in the Media Storage (flame-fms-metadata + flame-fms-storage) and then injected in the streaming server (flame-fms-streaming) to be delivered to the user.

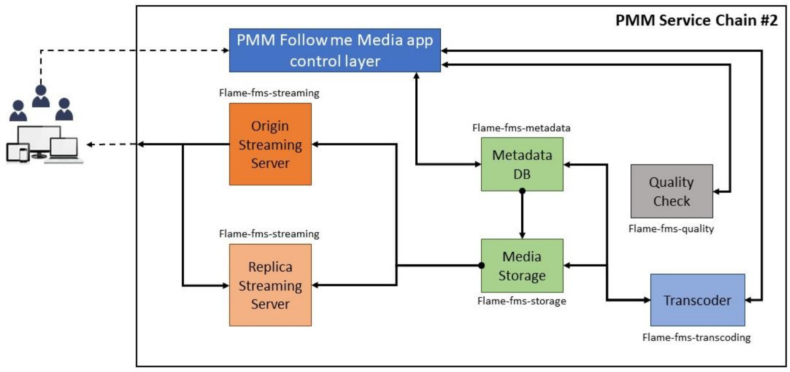

Figure 12: PMM service chain #2 media components interaction

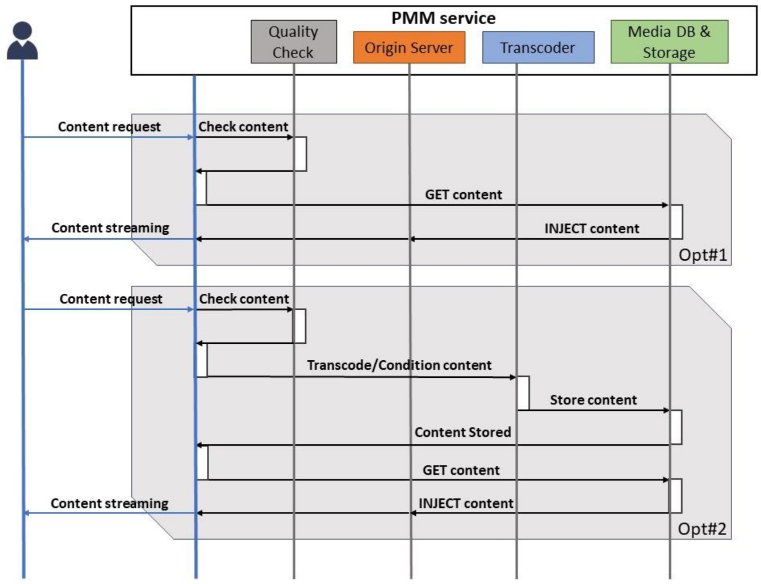

Figure 13: PMM service chain #2 simplified sequence diagram

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

4.4 INTEGRATION IN THE DESTINATION CITY

The PMM experimentation runs in the city area of Barcelona. Details about the infrastructure deployment are reported in Section 7.2. With respect to the provided infrastructure, which includes one core data centre located in the i2CAT premises, one edge cabinet in Calle Pere IV street plus several Wi-Fi access points placed on the street's lamps, the placement of the different media components for the proposed service chains is the following:

Service Chain #1: The PMM application, the Origin Server and the Media Storage (see Table 2 for details about infrastructure requirements) are deployed in the core data centre and interconnected between each other. During the adaptation of the service, according to the service's policies, specified through the service template, the Origin Server will be replicated/activated in the edge cabinet, in order to serve the user who is moving along Calle Pere IV. The Replica Server is connected to the same Media Storage located in the core data centre. Figure 14 depicts the mapping of the media service components with respect to the FLAME platform and Barcelona infrastructure.

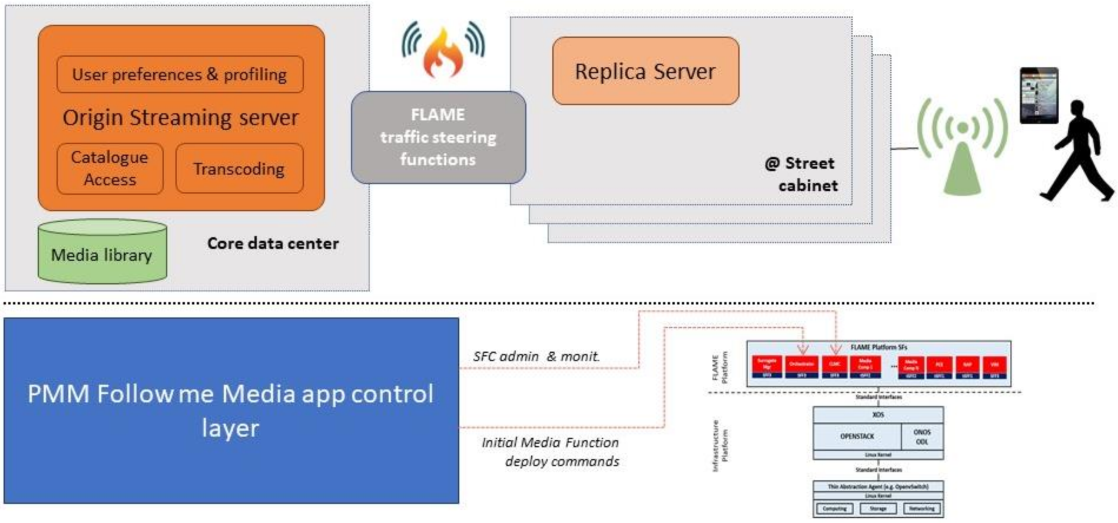

Figure 14: Service chain #1 media functions mapping with respect to the FLAME platform architecture

Service Chain #2: The PMM application and all the FMSs are deployed in the core data centre; the adaptation of the service is performed, according to policies, placing/connecting a replica of the Origin Server in the edge cabinet. As in the previous case, the streaming server at the edge will be connected to the same media storage and metadata DB connected to the one deployed in the core. Figure 15 reports the mapping between the service components in the service chain #2 and the FLAME platform with respect to the Barcelona infrastructure.

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

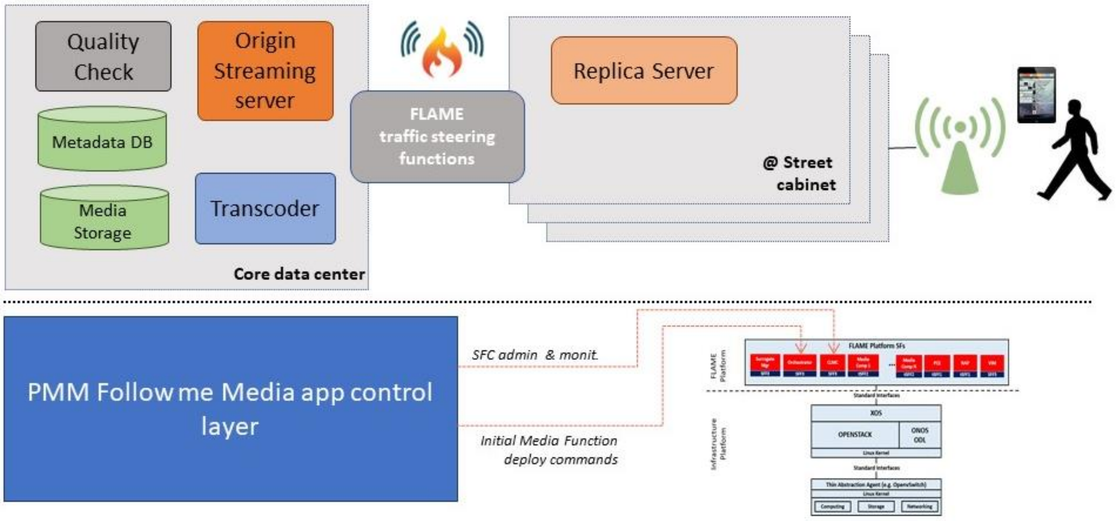

Figure 15: Service chain #2 media functions mapping with respect to the FLAME platform architecture

Table 2 PMM Media Components and related deployment flavours

Media Component	Deployment Flavour
Origin Streaming Server + Media Storage (NXW)	Virtual machine characteristics: ➔ Processors: 4vCPU Intel type, 2.4 GHz or better ➔ Memory: min 8GB RAM ➔ Storage: min 200 GB for central media library ➔ OS: at least Ubuntu 16.04 LTS Networks: 1 GbE for streaming + 1 Eth for server control

4.5 UPDATES AND CHANGES

With respect to deliverable D3.1, only one major change applies to the PMM experiment for the "My Screen follows me" scenario. In fact, due to limitations in CCTV hardware availability at the target infrastructure in Barcelona, the CCTV related live streaming of personal contents will not be experimented. However, the impact on the PIML characterization is minimal and limited to the use case u2.3 for the parts related to Interactivity on personal contents:

Interactivity

➔ ***u2.3*** – *The user must be enabled to browse his personal library of contents and start streaming stored media while on the move.*

In particular, the PMM experiment in use case u2.3 will not interact with live streaming devices like CCTV cameras and will be tested with more static media contents owned by the user and stored on the media server.

The other use cases identified for **Personalisation**, **Mobility** and **Localisation** remain invariant as per D3.1.

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

For what concerns targeted key stakeholders and value proposition in the PMM scenario, no updates or changes apply in this deliverable and the vision elaborated in D3.1 remains valid.

4.6 RELATION TO FMI VISION

All market analyses and forecast on the evolution of connected devices and major users' interests show a continued growth of the media applications use in mobile scenarios, in particular the video and home security ones streamed over the network. A recent publication of the "Visual Networking Index: Forecast and Trends, 2017–2022" by Cisco [18] confirms that globally the amount of VoD traffic by 2022 will be equivalent to 10 billion DVDs per month and 3% of all Internet video traffic will be due to video surveillance by 2022, up from 2% in 2017. We are experiencing a "Cord-Cutting" effect, i.e. there is an increasing migration from traditional and subscription-based fruition of media contents to more personal and mobile type of video services.

In this context, the Personalized Media Mobility scenario perfectly fits into the FMI vision outlined in previous Sec. 2, as also confirmed in [9] and [16], having its major value in the personalization aspects for users on the go.

From the media service provider point of view, the PMM service chain implemented for running the "My Screen follows me" scenario is designed for taking the maximum advantages from the capabilities offered by the FLAME platform. In the PMM scenario, the distribution of personalized media contents is optimized by integrating the media service components comprises the PMM service with the FLAME CLMC system, then enabling the automated adaptation of the service on the base of collected monitoring data. Upon the definition of service lifecycle policies that take into consideration the metrics that the application's components are able to export, the FLAME platform actuates the automated optimization/adjustment of the service. Indeed, the detection of the service degradation is a key feature towards the implementation of an efficient media service delivery that can adapt depending on the users' movement and preferences. In the PMM scenario specific case, the adaptation/optimization performed automatically by the FLAME platform consists in placing and connecting a new streaming server instance at the edge and closest to the end user, in order to serve the content maintaining or enhancing the QoE with respect to the initial setup of the service.

Progressing on this path, the current implementation of the PMM scenario addresses the PIML characterizations (as reported in D3.1 and updated within Section 4.5) by offering to the user the following functionalities:

- **Personalization**: the user is able to select the device through which he/she would like to stream the content. The streaming server is attached to a personal media library, then the user can stream personal content on the go, not only at his/her own premises.
- **Interactivity**: the user can manage his/her personal library from any device, for instance updating contents information.
- **Mobility**: the FLAME platform adapts the service chain while the user is moving in the area covered by the FLAME infrastructure (Barcelona infrastructure).

The development of PMM related scenarios and visions in FLAME allowed us to better understand the technical implications of developing such a service in a city infrastructure, the related PIML aspects of major interest for end-users and the potential positioning of the various stakeholders (media service providers, technology providers, platform and infrastructure providers, …).

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

5 SCENARIO: COLLABORATIVE INTERACTIVE TRANSMEDIA NARRATIVES

5.1 INTRODUCTION

The aim of this validation experiment is to tell interactive stories through AR within a city environment. The user enters the city, connects to the FLAME network, opens the app and can choose from a variety of stories that can be experienced in this region. Such stories can be used as a means of city tour guide, educational purposes or simply for entertainment. By augmenting the world, experienced through the smartphone camera, with 3D objects, a new level of immersion is reached and could bring historic scenes back to life in front of the user's eyes.

The smart city network allows to store 3D objects required for the stories in databases throughout the city – therefore, assets do not have to be pre-downloaded onto the device but can be streamed gradually as the story unfolds. To increase download speed, assets can be distributed in the city network onto several nodes, close to where they are required in the story. The scalable nature of the FLAME platform allows us to cope with high user activity, by automatically starting new databases throughout the city to distribute the load over more instances when latency is reaching critical levels at certain regions.

Related to the FMI Vision, our scenario utilizes the local storage capabilities and on-edge computation to increase the QoE in terms of reduced latency. Also, shared local experiences benefit from this architecture, as the availability of local servers allows almost zero latency state updates over all connected devices.

5.2 SCENARIO

5.2.1 The Story

The story told in our validation scenario is taking place in Bristol at the brink of the early 18th century. It is the age of pirates and overseas trading. While the age of discovery has ended, the marvels of the newly discovered world are still anchored deep within in the minds of the citizens and they are eager to talk about it. The user is meeting his old friend, Harry, a trader who was just recently struck by an ill fate: his ships coming from Portugal were raided by pirates. His fortune is gone. His ships are gone. But not his hope. Rumours reached his ears about a treasure buried in Bristol… and of course, he asks you, the user, for help in this quest.

The story will convey information about Bristol's history and some of its most famous inhabitants: Sir William Penn, a British admiral whose son founded Pennsylvania, Mr. William Tyndale, the scholar who translated the Bible into English, and several other personalities including the infamous pirate Blackbeard, whose made-up daughter Elizabeth will play a role in the story

5.2.2 Test Scenarios

We perform three different trials, gradually increasing the complexity of the scenario.

- **Pre-trials:** Test the intended SFC with 5 users. Push assets to databases and keep track of user progress in the story. Make sure assets can be downloaded and visualized in AR.

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

- **Small-trials:** Run a complete interactive story in AR with around 10 selected users and evaluate their experience in terms of responsiveness of the application and overall enjoyment.
- **Large-trials:** Run a complete interactive story in AR with around 25 users and evaluate their experience in terms of responsiveness of the application and overall enjoyment.

5.2.3 Service Function Chain

The service function chain utilized in our scenario is visualized in Figure 16. It consists of the Content Ingest & Storage FMS from Atos and a first version of the Story Compute SF. Each cluster is assigned a unique address to communicate with its contained Content Ingest & Storage FMS.

The **Story Compute SF** has the following tasks:

- Manage users logging in with the application
- Deliver the stories to users upon request
- Keep track of the users' progress in the story
- Keep track of the users' location in the city
- Determine which cluster the 3D assets have to be pushed to in order to enable a seamless story experience for the user.

The **Content Ingest & Storage FMS** has the following tasks:

- Store 3D assets in a local database
- Deliver 3D asset to user upon request

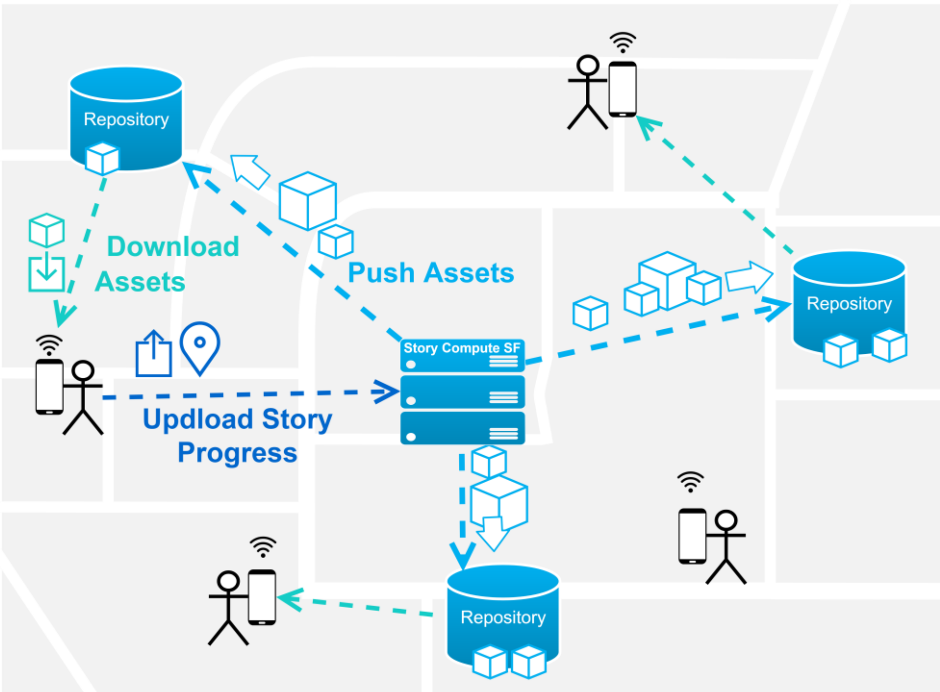

Figure 16 Service Function Chain Overview. Each Repository is an instance of the Content Ingest & Storage FMS.

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

5.3 MEDIA SERVICE USE CASES

The following main use cases result from the service function chain described above:

- **User Logs in:** Client sends POST request to Story Compute SF, which registers the name of the new user and returns a token with which the user can identify themselves in the future.

- **User Requests Story:** Client sends POST request with user token and story ID to the Story Compute SF, which returns the story file, containing all the required information for the client device to execute the interactive story.

- **User Starts Story / User Progresses in Story:** Client sends a POST request to the Story Compute SF with all the required information to locate the user within the story graph. The Story Compute SF processes all recent updates at regular intervals and bases its decisions on where to place assets within the city on this information (explained in the next paragraph). Additionally, it determines which assets the user already downloaded and stores this information in the user profile.

- **Determine Asset Placement** The main loop of the Story Compute SF runs through the most recent progress updates of each user. Based on these updates, the Story Compute SF determines which parts of the story this user could reach soon and prepares the assets required for these parts at the respective locations in the city - but only if the user has not already downloaded them. The Story Compute SF keeps metadata for each placed asset, indicating which user still requires the asset and for which part of the story. If the user no longer requires it or that part of the story is not reachable in the near future, this metadata is updated. If the metadata suggests that this asset is no longer required by any user, it is marked as removable.

- **Place Asset at Cluster:** StoryCompute pushes an asset to the desired cluster. If there is no repository instance running at that cluster, StoryCompute will send a specific trigger, such that an instance of the Content Ingest & Storage FMS is spun up at the required location. Once the new SF is running, StoryCompute can start pushing the assets.

- **User Requests Asset:** The client device sends a request to the Content Ingest & Storage FMS, which returns the specified asset. If it cannot find the asset in its storage, then other Content Ingest & Storage FMS instances within the SFC are queried until the required asset is found.

5.3.1 Future Changes

The Content Ingest & Storage FMS will eventually be replaced by a more elaborate vCDN by Atos, which slightly changes the way the content is placed on the clusters. However, since the vCDN is still in the design phase at the time of writing, we cannot elaborate on this in much detail. The major change is that the Story Compute SF will no longer have to push assets, but the different content repository instances will pull the required assets from each other. Each content repository instance will query the Story Compute SF at regular intervals about which assets it needs to fetch. The content repository instances are synchronized in such a way that requests for the same asset are sent by all instances requiring it at roughly the same time. This has the advantage that the GET requests can be multi-casted by the FLAME network, such that it only results in one single response, which is then delivered to all the content repository instances requesting it. This may also require that there is a master database, which contains all assets initially.

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

5.4 INTEGRATION IN THE DESTINATION CITY

The validation experiments will take place on Millennium Square in Bristol, where some of aforementioned famous Bristolians are present in form of life-sized bronze statues.

The story will be acted out on six different virtual stages. These 3D stages are mapped to specific locations on the Millennium Square with the help of GPS locations and AR markers. When the user gets into proximity of one of these locations, the stage is downloaded and placed as soon as the corresponding AR marker is found. The AR markers are used to properly orient the stage relative to the real environment. The respective stage-locations are indicated in Figure 17.

There are five cluster locations. For our validation experiment, all of them are used. The cluster location on “We the Curious” is used for the Story Compute SF (and the master content repository). All other cluster locations are utilized for asset content repositories with at least one instance of the Content Ingest & Storage FMS running on each one of them. Each of the 4 clusters is supposed to serve a different part of the square. However, their signal is likely to be spread over the whole square. Therefore, we have to ensure that the mobile devices used for testing are always connected to the closest WiFi, respectively the one with the strongest signal. This is one of the issues that arose from the change from a city-wide network to a square-wide network for testing.

Figure 17 The locations for the 3D stages. (Imagery @2018 Google, MapData @2018 Google)

5.5 UPDATES AND CHANGES

Updates and changes since D3.1 [2] are described here.

5.5.1 Location Change

Moving from a city-wide test network to one that only spreads over the area of Millennium square, triggered some issues:

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

- **More than one access point at any location**: Due to the small area, mobile devices can connect to any of the five access points from any place on the square. This makes the preloading of assets more difficult, even though we know at which location of the square they would be required. A solution is to install on each mobile device a service that ensures that it always connects to the strongest network (which is most likely also the closest one).

- **No branching structure**: Within city there are many side alleys and cross roads, where decisions have to be made. On a square, this structure is lost. This reduces the sense of agency within the story when it comes to deciding on where to go next.

5.5.2 Updates to PIML Characterization

Two of the PIML characteristics have slightly changed:

- **Personalization:** We build a user model to track what the user has already experienced. However, for the initial tests, the user model is not considered for making personalized suggestions of where to go next.

- **Interaction:** Interaction does not only happen through movement in the city but also through touch inputs on the device's screen to interact with the virtual characters and objects through the mobile device.

5.5.3 Updates to Media Service Use Cases

One use case described in D3.1 needs to be softened:

- **u4.6:** For the initial version of the Collaborative Interactive Narrative Experiences application, the use case of creating a recording of the whole experience is out of scope. It however is still a contender for future additions.

5.6 RELATION TO FMI VISION

With the upcoming wide-scale availability of 5G networks, more and more data can be streamed continuously onto mobile devices without the need to permanently store them. This allows for smaller sized applications that load their content as needed and discard it again when the application is closed. For the market of digital interactive experiences, this allows to offer different high quality 2D/3D content based on time of year (or day) or based on the device location.

Our application can leverage this development by providing different stories with possibly different actors and content at various geo locations. The FLAME network allows to spatially distribute content within a city, such that it is available near the location where it will be consumed. For a media service provider, this means reduced latency for downloading content onto the mobile devices of their customers. A smart content distribution algorithm could also minimize the server resources used for maintaining the application. With fewer users, the need for server resources is reduced and some of the clusters could be shut down and their content placed somewhere else. This is a feature that is already supported by the FLAME platform.

It is also possible that the assets would not even need to be downloaded to the mobile devices. This is if servers are equipped with strong GPU infrastructure, such that the experience can be rendered on the server and only the 3D overlay is delivered frame by frame to the client mobile device. The only things the server would require for this are the intrinsic parameters of the device's virtual camera and

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

the state of the story of each user. This would allow users with less powerful devices to also enjoy high-end graphics.

Having computational resources right on the edge also allows to synchronize the states of several devices with almost zero latency. This would for example allow a family to experience the same story, where the interactions of all family members would influence the progress of their shared story experience. Other examples are large-scale shared local experiences, such as a mini-game on a square with many participants in a local area, which requires synchronizing the devices of all participants before each user continues in their own story.

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

6 SCENARIO: AUGMENTED REALITY LOCATION BASED GAMING

6.1 INTRODUCTION

More than solely a visual and narrative experience, video games provide to the player a strong agency to explore and influence a whole environment. They create a unique setting where one can discover, experiment, try, fail, learn, and improve. The Game Technology Center at ETH Zurich is using augmented reality and localisation capabilities to bring these benefits to the physical world, expanding the action space of video games to the scale of a city.

As part of the FMI vision, city-wide video games could also be transformed to provide unique and personalized content for all the players. We demonstrate this possibility as part of the scenario for the game "Gnome Trader".

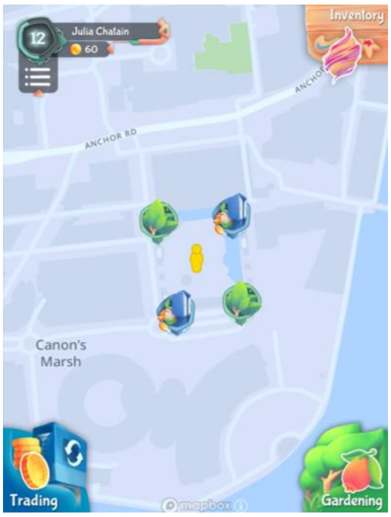

Figure 18 Example screenshot of the game "Gnome Trader"

6.2 SCENARIO

FLAME allows us to enhance video games with a mobile, location-based dimension that enables novel, urban game scenarios. Gnome Trader is an augmented reality, location-based video game where players move around in a designated space to interact with various game elements. The game brings two main mechanisms:

- At specific locations, the players find physical markers, which, when viewed through the phone's camera, become shops where gnomes are willing to sell or buy seeds. As the players trade these seeds with various gnomes, they affect the economy of the game world: the prices of rare seeds increase while the prices of more common seeds decrease.
- The players can create a garden anywhere they choose, plant the traded seeds, and make them grow to collect rewards. This way, players conquer their environment through novel interactions and expand the game world as they explore the urban space.

This scenario has several goals, expressed by the different game mechanics built in Gnome Trader. As a location-based game concept involving augmented reality, a strong goal of this scenario is to encourage real-world interactions, and, in particular, evaluate how players physically engage with their environment.

6.3 MEDIA SERVICE USE CASES

Gnome Trader employs FLAME's storage FMS. It uses the storage media service to store game assets, such as 3D models, textures, and sounds. The game assets are served to the player's mobile client in a localized way. That is, only game objects that relate to the player's current physical location and therefore are likely to be encountered in the current game session are downloaded and stored on the player's device. For Gnome Trader, this mechanism is particularly useful to retrieve the plants present

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

in a garden, or to retrieve the seeds present in a shop before a player visits it. It allows game developers to provide a large number of game assets to the game, which are not necessarily needed by every player and therefore would use unnecessary storage space on the player's device.

6.4 INTEGRATION IN THE DESTINATION CITY

As Gnome Trader is a location-based mobile game, the player's experience depends strongly on how the game is embedded into a physical space, such as Bristol's Millennium Square. Two main aspects need to be considered for a successful integration:

- The physical **locations** of the Gnome Trader shops, where the players can trade seeds, should be contained in the coverage of the FLAME network, be distant from each other, and cover all the space of the network.

- The **fertile areas**, where the players can grow plants (forming gardens) and collect rewards, can be planted anywhere in the coverage of the FLAME network on Millennium Square.

An example setup is depicted in Figure 19. Blue dots represent gnome shops and coloured shapes define different fertile areas at the Millennium Square.

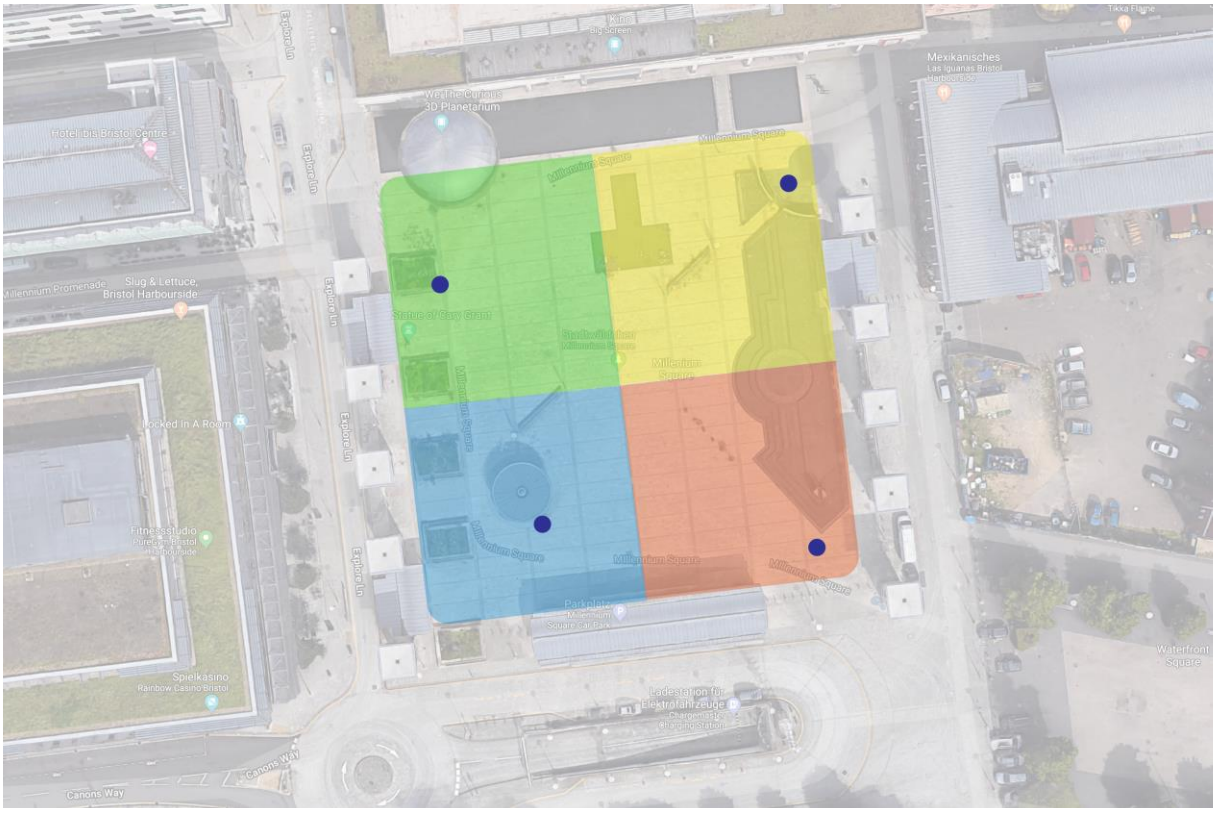

Figure 19 Example of fertile areas and shop locations in Millennium Square.

6.5 UPDATES AND CHANGES

Since D3.1 [2], the Gnome Trader scenario evolved to introduce "gardens", a game mechanic making use of the media services provided by FLAME to increase the agency the players have in the game

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

world. This is not only interesting to the player, but also from a game developer point of view, as gardens provide interesting data regarding how the players can conquer and extend their gaming world to fit their physical environments.

Additionally, along with the restructuring of FLAME to target Bristol's Millennium Square instead of the entire city, the Gnome Trader game world is also limited to the Millennium Square.

6.6 RELATION TO FMI VISION

In Gnome Trader, one of the main game mechanics relies on availability of unique trees at any location in space. In order to respect the game flow and provide the best experience to the players, these trees have to be downloaded within a short time frame. However, a player will only be willing to access the trees in their inventories and in nearby shops and gardens. Thanks to FLAME, we can ensure that these specific trees are available close to the user and thus can be downloaded much faster.

At this point of the project, the trees are generated by the game content creator and uploaded to FLAME so that the players can access them. In the future of the project, it could be possible to generate the trees directly on the platform as the players request their creation in game.

Finally, the trading mechanic of the game requires prices in the shops to be updated at a certain frequency, and to be consistent for every player trying to trade with a specific shop. In a future version, this update could be done at the access point closest to the shop.

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

7 CITY INFRASTRUCTURE

7.1 FLAME INFRASTRUCTURE IN BRISTOL

FLAME infrastructure in Bristol is part of the 5GUK Test Network that is deployed at Bristol city centre. In particular, at Millennium Square (MS) and We The Curious (WTC) Museum. This section provides an overview of the overall infrastructure available at MS, WTC and the Bristol-FLAME platform.

7.1.1 University of Bristol's 5G Infrastructure

In order to explore and validate the deployment of FLAME platform in an architecture that combines existing technologies and innovations, the University of Bristol has deployed a rich 5G network architecture comprised of several networking and computing technologies, interconnecting a significant area in the Bristol city centre. This overall infrastructure aims to provide a managed platform for the development and testing of new solutions delivering reliable and high-capacity services to several applications and vertical sectors here referred to as FLAME.

The University of Bristol's 5G network architecture, called 5GUK Test Network is a multi-site network connected through a 10km fibre with several active switching nodes, which are depicted in Figure 20. The core network is located at the Smart Internet laboratory at the University of Bristol and extra edge computing nodes are available in Millennium Square. As shown in Figure 21, the access technologies are located in two different areas in the city centre: Millennium Square for outdoor coverage and "We The Curious" science museum for indoor coverage. For the FLAME Platform, only the WIFI access technology is available in MS. The WIFI access points are located in four towers in each corner of the square.

Figure 20 Distribution of the 5GUK Test Network access technologies available for FLAME platform.

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

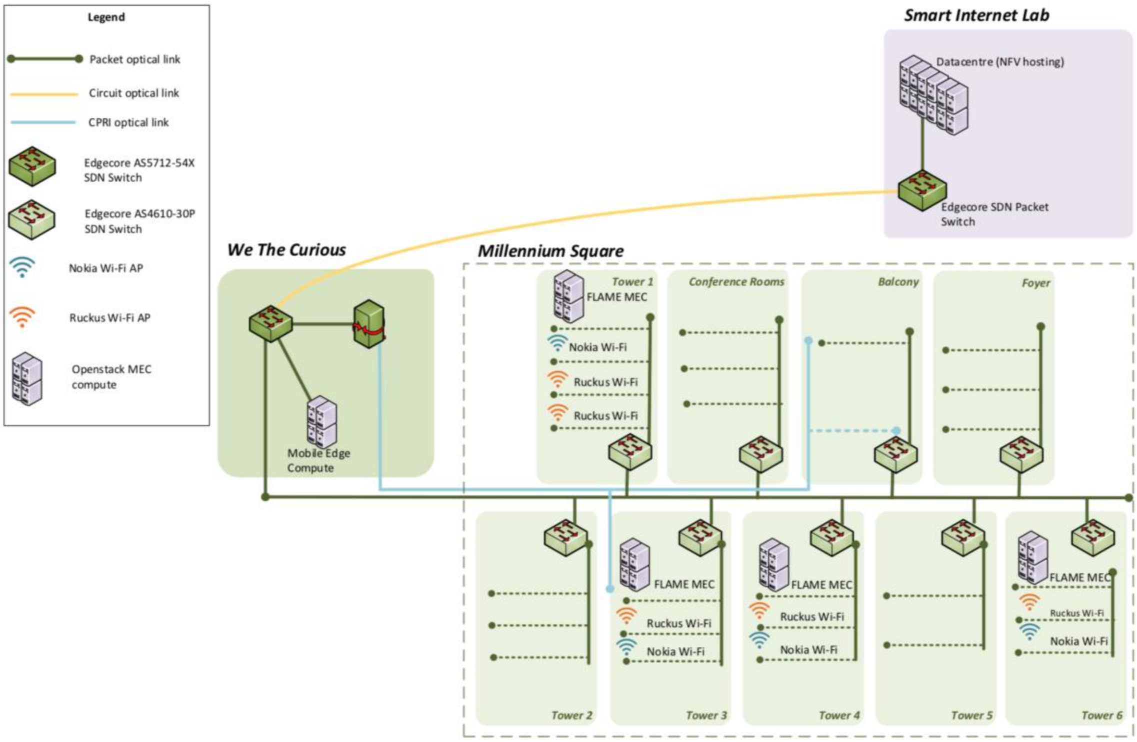

Figure 21 University of Bristol top level system architecture.

The aforementioned structure is used to support different vertical demonstrators, such as entertainment, finance, manufacturing and automotive testing.

The diverse range of access technologies are interconnected in sharing the same underlying system while being used by FLAME ecosystem to provide connectivity for the demonstrators, showcasing seamless integration between heterogeneous network components, an important concept in 5G.

The state-of-the-art radio access technologies deployed in Millennium Square will deliver high-bandwidth, high-bitrate and high-reliability connections to the user equipment, therefore enabling the usage of the network-intensive distributed applications developed by future FLAME experimenters if needed.

The SDN capabilities expressed by the NetOS controller will facilitate network slicing through optical, electrical and radio technologies via on-demand SSID creation, demonstrating another key concept in the 5G architecture that will be explored by FLAME to provide a multi-tenant environment, where the multiple demonstrators, or even final users, can coexist independently with different connectivity specifications.

Finally, the University of Bristol 5G infrastructure delivers an automated and programmable environment, that will be used by FLAME southbound interface to create fully integrated orchestration for both application components and network services.

7.1.2 FLAME Platform

Figure 22 shows the logical FLAME platform architecture deployed at Millennium Square in Bristol. A set of four towers has been allocated to host FLAME Mobile Edge Computing nodes. Each tower has a compute node based on OpenStack Ocata NOVA. Each compute node is connected to an EdgeCore SDN switch that is connected to a single SDN switch located at WTC.

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

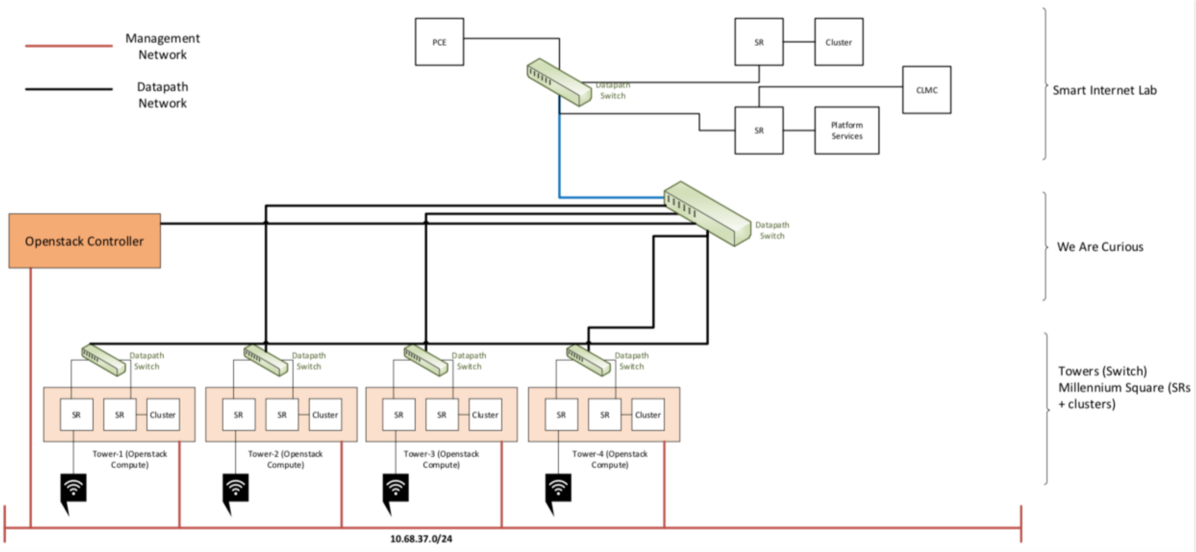

Figure 22 Bristol-FLAME platform architecture

The four compute nodes and the edge core SDN switches are connected to the SDN controller based on FloodLight, also located at the WTC. One compute node is deployed for each of the four towers. The Bristol infrastructure constituent equipment and capabilities include the following:

- Multi-vendor software-defined networking (SDN) enabled packet switched network
  - Corsa switch (Corsa DP2100)
  - Edgecore switch (Edgecore AS4610 series & AS5712-54X)

- SDN enabled optical (Fibre) switched network
  - Polatis Series 6000 Optical Circuit Switch.

- Multi-vendor Wi-Fi
  - SDN enabled Ruckus Wi-Fi (T710 and R720)
  - Nokia Wi-Fi (AC400) (Not Available for FLAME)

- Mobile Edge Computing node
  - Dell Power Edge R430 Server
  - Each MEC has available to the experimenter: cores = 15, RAM = 29GB, disk = 900GB

7.2 FLAME INFRASTRUCTURE IN BARCELONA

This section details how the Barcelona infrastructure was designed to provide a space for experimenters to play with features that will be typical in future 5G deployments, such as edge computing and wireless backhauling that are key for high performance processing and flexible infrastructure deployment. In Barcelona, the infrastructure that hosts the FLAME architecture can be split into two major segments that are hosted in different areas of the city. First, the Main DC that features powerful computing equipment and enough resources to host the core elements of the FLAME platform is located at i2CAT in the *Zona Universitaria* (see Figure 23). Second, the on-street deployment that includes the street cabinet server and the lamp posts hosting the RAN elements is located in the so called *22@* area. Furthermore, this area also has a dedicated space at IMI that hosts networking equipment that acts as relay for the traffic between the main DC and the on-street deployment. This connectivity is provided over a 10 Gbps fibre link with redundancy.

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

Figure 23: FLAME is deployed in two main locations, the Zona Universitaria (orange circle) hosting the main DC and the 22@ area (blue circle) hosting on-street deployment including the street cabinet and lamp posts. Both locations are connected over 10 Gbps fibre.

The *22@* location can be split into two spots with FLAME infrastructure presence: i) the IMI premises (Glories Office), where the networking equipment that enables the connection between main DC and on-street equipment is hosted and ii) the Pere IV installations, i.e. the lamp posts that are equipped with Wi-Fi nodes for the RAN, as well as the street cabinet that hosts the cabinet server and networking equipment (see Figure 24) required for high performance edge computing. The latter is required to interconnect the on-street elements with each other and with the IMI premises. Figure 25 shows all the elements of the infrastructure deployment used to host FLAME in Barcelona.

The Main DC located in the Omega building is composed of three servers which are configured to provide high availability and failover replication. These servers logically work as one server, primarily containing Virtual Infrastructure Manager Openstack services. Thus, the Main DC provides an Openstack Controller (including different internal services for networking and image storage) and two compute nodes, named “ias1” and “ias2”, to hold VM instances in the Main DC. Furthermore, the hardware of these servers is composed of:

- 4 x 3GHz CPU, 32 GB RAM, 240 GB SSD. IAS0 server performs a controller service
- 6 x 2.4GHz CPU, 96 GB RAM, 1900 GB SSD. IAS1 server performs controller, compute and volume services
- 6 x 3.5GHz CPU, 96 GB RAM, 1900 GB SSD. IAS2 server performs controller, compute and volume services

One of the servers, IAS1, is responsible for hosting the SDN Controller of the infrastructure, which has been switched from OpenDaylight to Floodlight in order to assure compatibility with the NAP software elements. This SDN Controller will serve both the FLAME platform and the wireless backhaul.

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

Figure 24: Lamp post with Wi-Fi node and directive antenna (left) and street cabinet with networking equipment (right)

In Pere IV a total of 4 lamp posts have been selected to host RAN nodes to offer high throughput Wi-Fi connectivity to the experimenters. The location of the lamp posts is shown in Figure 26, along with the position of the street cabinet, where the cabinet server and networking equipment are stored. Each lamp post is equipped with what we call a wireless node that hosts a Wi-Fi transceiver on top of which Wi-Fi access points (APs) can be instantiated to provide connectivity to the on-street users. These wireless transceivers support the IEEE 802.11ac standard, operating in the 5 GHz unlicensed band and they offer stable transmission speeds of around 200 Mbit/s, with peak throughputs of up to 300 Mbit/s for single users close to the lamp post (assuming that the user equipment supports this standard and at least 2x2 multiple-input multiple-output). Thus, the theoretical aggregated throughput for the RAN amounts to over 1 Gbit/s, enough capacity to sever high quality media services in the deployment area.

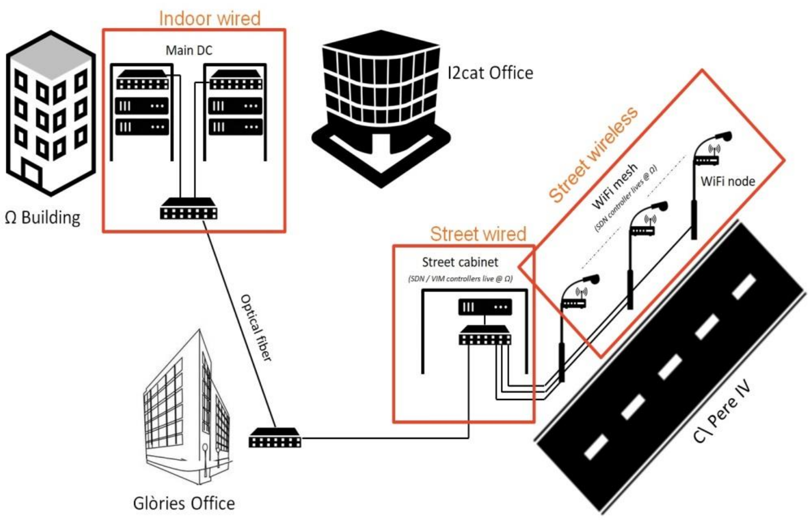

Figure 25: Architecture of the FLAME Barcelona infrastructure

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

Further, the wireless nodes are equipped with one or two *wireless backhauling* transceivers, again supporting the IEEE 802.11ac standard. For high performance and link stability, directive antennas were installed on the lamp posts in such a way that they would each point towards its neighbouring lamp post(s). As can be seen in Figure 26, the lamp posts have been chosen in such a way that Line-of-sight (LoS) is guaranteed between neighbours. This is important for stability, but also for performance. The wireless backhaul in the on-street deployment is required, since from an experimenter's point of view only one of the lamp posts is actually connected to the wired backhaul (towards the cabinet/Main DC/Internet). As such, it is necessary to carry any traffic from or to any of the other lamp posts over the wireless mesh backhaul towards this "gateway" lamp post.

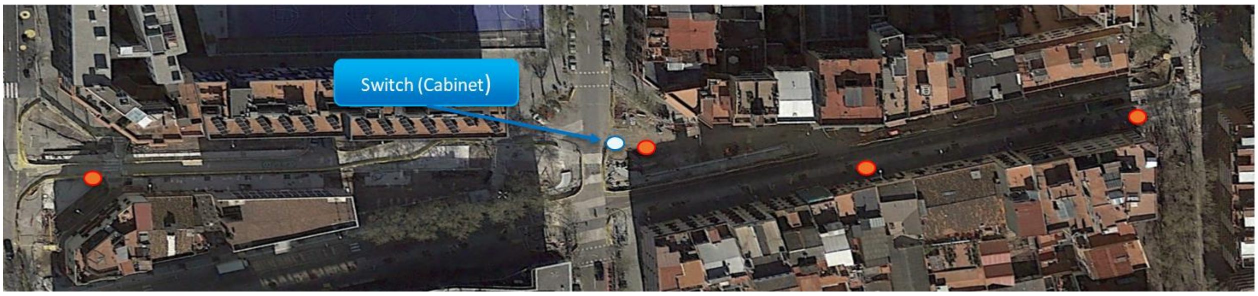

Figure 26: Location of the cabinet and the four access points mounted on street lamps.

The validation of link stability and performance yielded the results shown in Figure 27. The lowest average speed observed between any of the lamp posts is 180 Mbit/s. It shall be noted here that the FLAME platform elements, i.e., software elements deployed in the Barcelona infrastructure to provide the FLAME services, require a certain amount of computing resources. In order to host the edge services in the cabinet server and to still be able to offer enough compute resources to experimenters to host their services, the maximum number of lamp posts that can be integrated in the FLAME platform is limited (as each additional lamp post requires an additional FLAME software element along with some computing resources).

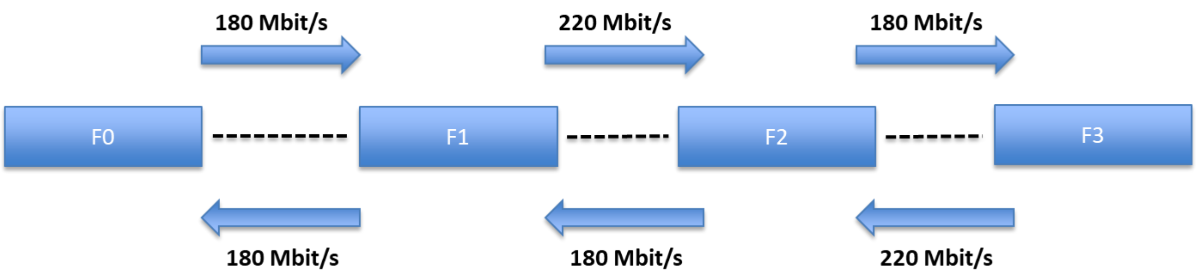

Figure 27: Experimental evaluation of the backhaul link performance in Pere IV street. The nodes are labelled from F0 to F3, corresponding to the leftmost node (F0) in Figure 26 to the rightmost one (F3).

In order for FLAME to route traffic towards specific lamp posts, FLAME's SDN fabric needs to be able to configure routing rules in the lamp posts that allow the communication between media services, users (Wi-Fi clients) and any other service instantiated and managed by FLAME. For that, the wireless nodes host OpenvSwitch instances that include all the ports of the wireless backhaul and the wired backhaul ports, allowing traffic to be routed between the lamp posts and the cabinet/main DC. The OvS need to be configured with FLIPS rules injected by an SDN controller that forms part of the infrastructure. For the Barcelona deployment, a Floodlight instance is used as SDN controller in a virtual machine governed by Barcelona's OpenStack. Figure 28 shows the topology with 3 lamp posts as seen by the Floodlight SDN controller.

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

For more details, please refer to D5.1 [17], which includes further diagrams and hardware specifications.

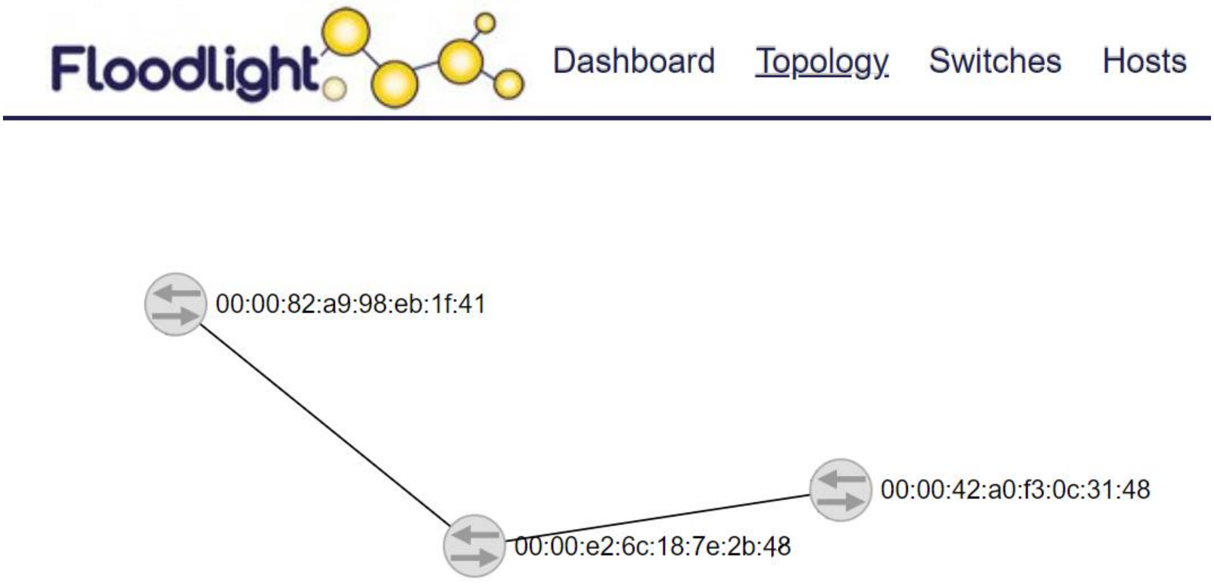

Figure 28: The Floodlight SDN controller's view of the virtual switches running on three lamp posts

Figure 29 shows a diagram of the connectivity and VLAN setup of the infrastructure. Main DC (Omega) extends a set Layer 2 networks using VLANs (which range from 0.10 to 0.500) focused to enable services along the entire network architecture, to the small Cabinet DC and the wireless backhaul. Those services highlight the management network, wide area network, data network, SDN control network and secure inbound access network. In addition, another network can be used to manage the infrastructure internally. On the edge (Cabinet DC and wireless backhaul), we can find the access network where users can connect to the different lamppost nodes. The edge is designed to simulate that each lamppost holds its own virtual infrastructure. We achieve this by using these VLANS and focusing all the virtual infrastructure resources in a single point, the Cabinet DC. Each lamppost can dispose of its dedicated NAP. Since the whole network – from the Main DC to the very edge- is connected over a L2 network, the delays between user devices and the services running in the edge or even in the main DC are very low, which is an enabler for FMIs with low latency requirements, such as the applications of the different scenarios presented in previous sections.

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

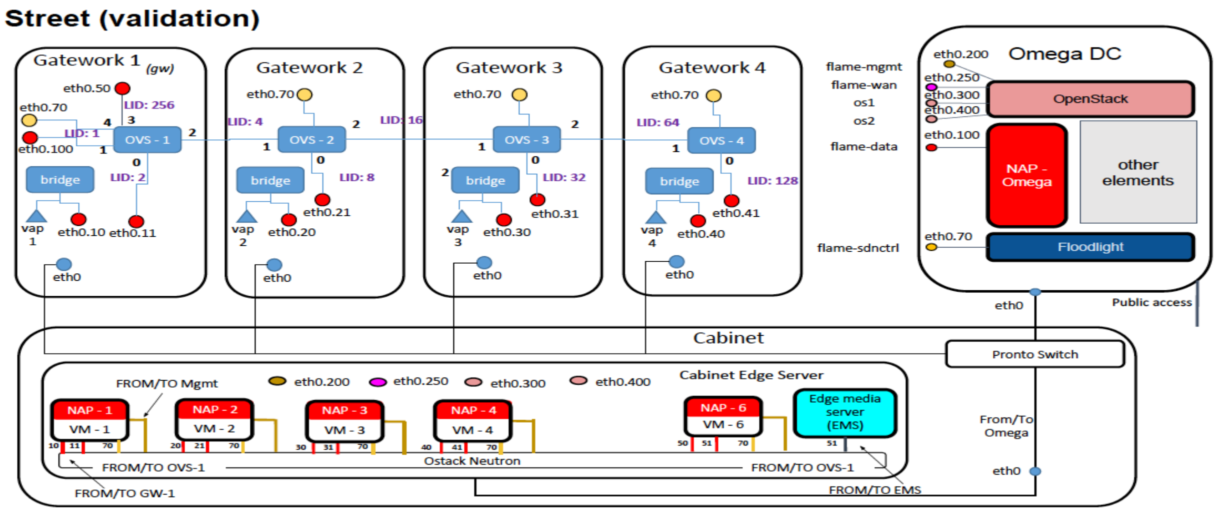

Figure 29 Connectivity and VLAN diagram Barcelona Testbed

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

8 CONCLUSIONS

This deliverable has presented the second iteration of the developing vision of FLAME within the Future Media Internet (FMI). The FLAME FMI vision envisages the landscape of media services in the next few years and considers new technological paradigms and principles that inspire 5G. The vision was presented in the context of general user requirements and specific use cases that consider 5G capabilities, which are demonstrated by the FLAME platform. The vision also considers the three principal stakeholders: Media service providers, platform providers, and infrastructure providers & owners. These three stakeholders are represented by the FLAME consortium

The deliverable has also provided an updated description of each of the four initial FLAME validation scenarios. These scenarios are representative of broadcast, gaming or transmedia vertical markets. Each scenario was introduced by describing its relation to the FMI vision. The scenarios and media service use cases were described, as well as details about the integration in the destination city. Updates and changes since the first version of the deliverable, D3.1 [2], were highlighted. Each scenario was concluded by detailing its relation to the FMI vision.

The FMI vision realized by the two city infrastructure locations was also described. The infrastructure locations are designed according to principles required to demonstrate key concepts of 5G architectures with the FLAME platform. The capabilities offered by the infrastructure provide important constraints on the validation scenarios and reflect the importance of considering infrastructure when developing the FLAME platform and media service scenarios.

The evaluation of the four validation scenarios will be presented in D5.5 "Insights from Broadcast, Gaming and Transmedia Experiments". These insights are expected to validate a subset of the FLAME FMI vision.

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

REFERENCES

[1] Boniface, M., Trossen, D., and Calisti, M. "Tackling User-Centric Media Demands through Software Defined Infrastructures." NEM Summit, Portugal. 23 – 25 Nov 2016. URL: https://eprints.soton.ac.uk/404229/1/404229.pdf

[2] D3.1 FMI Vision, Use Cases and Scenarios v1, available at https://ict-flame.eu/wp-content/uploads/sites/3/2017/10/D3.1-FMI-Vision-Use-Cases-and-Scenarios-v1.1.pdf

[3] D.3.3 FLAME Platform Architecture and Infrastructure Specification v1.0, available at https://ict-flame.eu/wp-content/uploads/sites/3/2017/10/D3.3-FLAME-Platform-Architecture-and-Infrastructure-Specification-v1.1.pdf

[4] European Broadcasting Union. "5G and Public Service Media (PSM): Opportunities in distribution of audiovisual content and services" Executive summary. February 2018.

[5] Atos. "Toward the interactive fan experience". Look Out 2020+ series. Industry Trends. Sports & Entertainment. June 2018. Available here: https://atos.net/content/mini-sites/look-out-2020/assets/pdf/ATOS_LOOK%20OUT_SPORTS.pdf

[6] Atos. "Accelerating the race to value". Look Out 2020+ series. Industry Trends. Telecom. June 2018. Available here: https://atos.net/content/mini-sites/look-out-2020/assets/pdf/ATOS_LOOK%20OUT_TELCO.pdf

[7] Ulrich Reimers. "The future role of broadcast in a world of wireless broadband". Keynote in 2015 IEEE Symposium on Broadband Multimedia Systems and Broadcasting. June 2015.

[8] FLAME Consortium. Deliverable D2.1 "Market Analysis for the FLAME FMI Ecosystem". June 2017.

[9] 5G-PPP and NEM. "5G and Media & Entertainment. Whitepaper". January 2016.

[10]De Bono, Edward. New Thinking For The New Millennium. Viking Adult, 1990.

[11] United Nations Announces Global Media Compact, UN.org, 23 September, 2018, Available here: https://www.un.org/sustainabledevelopment/blog/2018/09/media-compact-launch/. At 90 years old, European television is bolder than ever, UNRIC.org, 27 November 2018. Available here: https://www.unric.org/en/latest-un-buzz/31190-at-90-years-old-european-television-is-bolder-than-ever

[12]Mean Opinion Score, see also: https://en.wikipedia.org/wiki/Mean_opinion_score

[13]Casas, P., & Wassermann, S. (2018). Improving QoE prediction in mobile video through machine learning. Proceedings of the 2017 8th International Conference on the Network of the Future, NOF 2017, 2018–January (November), 1–7. https://doi.org/10.1109/NOF.2017.8251212

[14]Columbus Podcast Meetup presentation on demographics and psychographics by Gary Monti May 22, 2018. https://vimeo.com/271948810

[15]Is social targeting the future? By Richard Pentin. 01 July 2010. https://www.levidepoches.fr/contagiousideas/2010/07/is-social-targeting-the-future-by-richard-pentin.html

[16]van Eeden, E., & Chow, W. (2018). Perspectives from the Global Entertainment & Media Outlook 2018–2022 Trending: Trending now: convergence, connections and trust. Retrieved from www.pwc.com/outlook

[17]D5.1 FLAME Replication Process v1, available at https://ict-flame.eu/wp-content/uploads/sites/3/2018/02/D5.1-FLAME-Replication-Process.pdf

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

[18] Cisco Visual Networking Index: Forecast and Trends, 2017–2022, available at: https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white-paper-c11-741490.pdf

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