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

D5.5: Insights from Broadcast, Gaming and Transmedia Experiments

Steven Poulakos (DRZ), Manuel Braunschweiler (ETH), Fabio Zund (ETH), Marc Godon (VRT), Klaas Baert (VRT), Lorenzo Neri (NXW), Gino Carrozzo (NXW), Dirk Trossen (IDE), Stephen C. Phillips (ITINNOV), Michael Boniface (ITINNOV)

22/11/2019

This deliverable presents the insights from broadcast, gaming and transmedia experiments resulting from the validation experiments. The experiments use and evaluate FLAME's methodology and platform capabilities through experiments deployed at FLAME infrastructures. Experiments involving human participants were conducted at the FLAME infrastructure locations in Bristol and Barcelona. The validation experiments not only validated the technical capabilities, but validated and improved the methodology, tools and documentation for delivering localized media services on FLAME.

Work package	WP 5
Task	Task 5.1 and 5.2
Due date	30/04/2019
Submission date	22/11/2019
Deliverable lead	DRZ
Version	1.0
Authors	Steven Poulakos (DRZ), Manuel Braunschweiler (ETH), Fabio Zund (ETH), Marc Godon (VRT), Klaas Baert (VRT), Lorenzo Neri (NXW), Gino Carrozzo (NXW), Dirk Trossen (IDE), Stephen C. Phillips (ITINNOV), Michael Boniface (ITINNOV)
Reviewers	Carlos Alberto Martin (ATOS) and Dirk Trossen (IDE)

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 presents the insights from broadcast, gaming and transmedia verticals resulting from the validation experiments. The experiments use and evaluate FLAME's methodology and platform capabilities through experiments deployed at FLAME infrastructures.

This deliverable includes the presentation of high-level insights developed by FLAME partners about new media services enabled by FLAME. The high-level insights summarise new media service trends and motivate the consideration of new media spaces that combine digital media with the physical environment. The concepts have inspired the validation experiments in the vertical media segments described below.

Participatory Media for Interactive Radio Communities: The first experiment, on Mobile Journalism, revealed the potential value of media creation and collaboration tools. Placing these tools on the FLAME edge network enables the journalist to be quickly responsive to creative needs. This experiment confirmed the value of a virtual director room to support the creative act of media production. The second experiment demonstrated the value of deploying interactive media experiences in a public place. The experiment demonstrated value with linking a digital social network with a physical city space.

Personalized Media Mobility in Urban Environments: Experiments evaluated how media service providers can serve users on the go within a smart city. The FLAME platform automatically instantiates content caches and adapts routing in media service chains to deliver the best streaming experience while on the move. The experiments demonstrated the ease of deployment offered by the FLAME platform as well as automatic load balancing offered by FLIPS-based service routing.

Collaborative Interactive Transmedia Narratives: The experiments on City-wide Storytelling demonstrate how the FLAME platform offers unique capabilities for intelligently and efficiently delivering multimedia assets within an urban environment. The use of location-based augmented reality provided a means to integrate digital content within physical spaces. The FLAME platform is demonstrated to deliver digital content from these physical spaces and can be orchestrated to provide these services only as required. The result not only adds value for the current application, but demonstrate the potential for incorporating additional media service, such as state synchronisation for shared media experiences.

Augmented Reality Location Based Gaming: The augmented reality location-based gaming experiments demonstrated that the FLAME platform is capable of delivering 3D assets quickly and efficiently for video game applications. The experiments specifically demonstrated the value of the opportunistic multicasting capabilities of the FLAME platform. Players were immersed in the virtual augmented reality environment and mostly unaware of the 3D asset loading in the background.

The four validation experiments have helped drive the development of the FLAME platform and processes. The validation experiments provided the first concrete requirements for the platform and informed the technical roadmaps for the CLMC and FLIPS components. The experiments also helped to inform the FMS roadmap, with several components initially tested by these validation experiments. The validation experiments not only validated the technical capabilities, but validated and improved the methodology, tools and documentation for delivering localized media services on FLAME.

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

ABBREVIATIONS

3GPP	3rd Generation Partnership Project
AR	Augmented Reality
CLMC	Cross Layer Management and Control
FMI	Future Media Internet
FMS	Foundational Media Service
HDR	High Dynamic Range
KPI	Key Performance Indicator
MEC	Multi-access Edge Computing (formerly known as Mobile Edge Computing)
MR	Mixed Reality
NFV	Network Function Virtualisation
OTT	Over The Top
PMM	Personalized Media Mobility
QoE	Quality of Experience
QoS	Quality of Service
SDN	Software Defined Networking
SFC	Service Function Chain
VR	Virtual Reality
WCG	Wide Colour Gamut

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

1 INTRODUCTION

This deliverable presents the insights from broadcast, gaming and transmedia experiments performed by the FLAME validation partners. The experiments use and evaluate FLAME's methodology and platform capabilities, involving human participants at the FLAME infrastructure locations in Bristol and Barcelona. Additional insights were informed by Quality of Service (QoS) experiments conducted in the FLAME Sandpit infrastructure hosted at ITINNOV in Southampton.

This deliverable includes the presentation of high-level insights developed by FLAME partners about new media services enabled by FLAME (see Section 2), those specific insights demonstrated by the validation experiments. The high-level insights summarise new media services and trends. The section also presents the importance of considering new media spaces as a combination of media and the utility it provides in the context of physical spaces. The concepts have inspired the validation experiments in the vertical media segments described below.

1.1 VALIDATION BY BROADCAST VERTICAL

Two pioneering FMI experiments were designed and implemented to explore the changing way that consumers participate and access broadcast media on the move.

The validation scenario "Participatory Media for Interactive Radio Communities" (see Section 3) explores the innovative technical capabilities and value of the FLAME platform for developing new media interaction inside the city. Experiments explored the potential of utilizing localized media services to support mobile journalism activities and to support interactive media experiences that link digital social networks with physical city spaces. The results inform public broadcasters and civic organizers about the potential symbiosis between local media interaction and city life.

The validation scenario "Personalized Media Mobility in Urban Environments" (see Section 4) evaluates how media service providers can serve users on the go within a smart city. The FLAME platform automatically instantiates content caches and adapts routing in media service chains to guarantee the best streaming experience while on the move. The experiments explore the personalized media mobility "my media follows me" service while users are on the go and for the sharing of personal media in aggregation areas. The results inform media service providers of interactive digital platforms and personal media management systems to seamlessly deliver content across physical locations.

1.2 VALIDATION BY GAMING AND TRANSMEDIA VERTICALS

Two more pioneering FMI use cases were implemented to explore interactive media services for storytelling and gaming.

The validation scenario "Collaborative Interactive Transmedia Narratives" (see Section 5) demonstrates how the FLAME platform supports the delivery of location-based augmented reality (AR) stories within an urban environment. The aim is to augment physical locations allowing mobile users to experience increasingly interactive and content rich spaces that bring together digital and physical worlds. This requires new and advanced media services that are localized within the physical environment. The FLAME platform enables intelligent orchestration and management of these media services. The results inform media services providers intending to spatially distribute interactive, transmedia content for contextually appropriate delivery at physical locations within a city.

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

The validation scenario "Augmented Reality Location-based Gaming" (see Section 6) presented a location-based urban game that uses FLAME to deliver content that transforms locations into an interactive game. The FLAME platform transforms simultaneous requests for the same asset into one single request, which is then sent to the 3D content database. The evaluation indicates that the FLAME platform can deliver 3D assets quickly and efficiently if the underlying code is tuned to benefit from the multicasting capabilities. The results inform media service providers aiming to deliver 3D game assets to all players throughout the game play.

1.3 ORGANIZATION OF THIS DELIVERABLE

Section 2 "New Media Services" provide insights informed by experimentation about the future of new media services enabled by FLAME. Sections 3 - 6 present the outcomes of experiments conducted at FLAME infrastructures delivering insights into technical performance, user acceptance and viability of these FMI use cases. These experiments demonstrate examples of enhanced media products and provide useful feedback to evaluate the FLAME platform.

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

2 NEW MEDIA SERVICES

This section presents identified FMI media service trends as well as the emerging concept of *media spaces*, which combine physical with content interaction. We then contrast those new media trends against the emerging 5G capabilities, realised through the FLAME platform, and how our validation experiments utilise those capabilities to showcase some of the identified new media trends.

2.1 NEW SERVICES AND MEDIA TRENDS

The current media panorama is characterised by the creation and deployment of **new media services** jointly with new user habits and consumption patterns. Some of the most significant **new services and media trends,** as described in [1], include the following categories:

1. Demand for **improved quality**, including **more pixels** (e.g. beyond 4K resolution) and also **better pixels** (e.g., HDR or WCG). This category integrates characteristics such as dynamic range, resolution, frame rate, spatial audio and video colour gamut. This trend is consistent with the new capabilities of cameras and acquisition equipment offering new capabilities to provide high quality media material. The trend is also consistent with the features of consumer electronics according to the manufacturing industry strategy. This media trend impacts the throughput required for media transmission.

2. New technologies that provide **an enriched user experience**, beyond the unidirectional presentation of video and audio. **Gaming, AR, VR and MR** are good examples of this trend. Users look for **more engaging, immersive and exciting experiences**. For example, **360-degree media**, enabling virtual reality (VR) and other immersive applications, is another example of improved experience that places new demands on media services.

3. **Localised contents and experiences** can be linked to a physical or symbolic space, such as a sports match or a culture festival in the city. This offers new opportunities to achieve a better enjoyment or improved utility from linking digital and physical information.

4. **Content consumers are becoming content producers**. Modern mobile equipment enables quality media acquisition and social networks enable the transmission in real time. This phenomenon is also connected to the localised experiences described in the previous bullet: many users desire to disseminate their participation in public events and social networks enable this desire instantaneously.

5. Users wish to watch **any content anywhere and at any time**, e.g., as they navigate across a city or during a daily commute.

2.2 NEW MEDIA SPACES

Media service providers are not simple OTT operators on FLAME-enabled facilities: due to the multilayer FLAME approach, the media provider not only uses the underlying network capabilities, but the network is adaptive itself to optimise the new services quality, according to the policies established by the media service provider. [1]

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

The city is not only the place where media services are distributed and consumed. The city is a place for experimentation, participation and interaction, which stimulates the contact among the citizens and which links the citizens with urban symbolic values, like the city history.

The public space even stimulates a sort of city gamification. Citizens repurpose their surroundings for playful behaviour: squares become the stage for a flash-mob, obstacles become a challenge for Parkour runners, and lines on the pavement become traps we should not step on. This kind of city-gamification can be partially shifted into the virtual world – especially within a smart city, where computational resources are readily available and data exchange between users is faster than ever. Augmented Reality (AR) technologies enable mobile devices to recognize city features, overlay them with fantastic 3D models, and allow virtual interactions with those objects. Therefore, AR is a key technology to repurpose parts of the real world into an environment of play. Smart city infrastructure can be used in a variety of ways: sensors, for example, can give input to the virtual representation and local servers synchronize the state of the virtual augmentation for several players, such that all of them can see the changes other people apply to the virtual world in real-time. A key factor for these capabilities is speed. Not only from the mobile devices but also the server infrastructure. Synchronization for several players must be instantaneous to give the impression of the augmentation being real.

The high-level insights and examples presented above have guided the development of the FLAME platform and new media services utilizing it. The following sections present specific insights demonstrated by the validation experiments.

2.3 MAPPING ONTO EMERGING 5G CAPABILITIES

The advance of **emerging 5G infrastructures** and the **programmability of those infrastructures** sees the trend towards increasingly deploying localised infrastructure capabilities in the form of so-called 'private network deployments' that even include localised radio spectrum for high bandwidth applications. Specifically, **local event sites**, such as squares (e.g., as in Bristol), cities (e.g., as in Barcelona) or buildings (e.g., at KCL in London) can therefore enter direct business relationships with service platform providers as well as media service providers for optimised localised experiences along the aforementioned trends for new media services.

The FLAME platform targets such private network deployments, enabling the deployment of these new media services localised on the city facilities, while enabling the **orchestration and management of media services,** optimising both compute and network resources. This is achieved through a distributed, virtualisation enabled platform capability that is being realized over the localised infrastructure, while aligning the core FLAME technologies with 5G standards and solutions, as outlined in [1].

In order to support with the aforementioned media trends, the FLAME platform allows for an **intuitive description of the network service** as service function chains (SFCs), offering the programmability, deployment and management of media services, while enabling a service DevOps process for faster new media service creation and deployment. The average '*service creation time*' is one of the main KPIs defined by the European 5G PPP initiative targeting a reduction '*from 90 hours to 90 minutes'*. Through its infrastructure replication toolchain, the FLAME platform itself can be deployed within this outlined KPI, while media services themselves are deployable in a few minutes or less.

The FLAME capabilities [10] optimise the performance of media services, including an innovative routing solution. The overall FLAME benefits for service deployment are: low latency compute and

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

delivery (key to new interactive media services at the far edge of the network), fast and dynamic service request routing (key to personalised and mobile services), multicast delivery of HTTP responses (key to scalability), network-level indirection (key to replication of content based on local relevance), and more secure content objects (key to content replication while preserving content security and end user privacy).

2.4 BRINGING MEDIA TRENDS AND 5G CAPABILITIES TOGETHER

The FLAME validation scenarios demonstrate how the aforementioned platform capabilities are impacting the creation and deployment of new media services. The FLAME validation scenarios also present many of the characteristics and trends of media services described above. Moreover, FLAME validation partners have identified additional use cases to be offered in the facilities, utilising FMS (or Foundation Media Services, developed in FLAME) components to build more complex media services.

Furthermore, FLAME has produced a whitepaper [3] that explores the intersections between 5G and Future Media Internet (FMI) from the FLAME contribution perspective, to outline the drivers shared by 5G and FLAME and the main contributions of FLAME to 5G developments. These capabilities of the FLAME platform to support new services and the set of FLAME benefits cannot be seen without considering the close relationship between FLAME and 5G deployments in the city facilities. FLAME is deploying the similar technological paradigms that are necessary for 5G to go for a massive deployment, such as software network virtualisation (NFV/SDN) and automated service deployments as virtualised components that make the most of the edge network too (MEC scenarios).

The existence of deployments in realistic city facilities is essential in order to test new services and to extract measures and conclusions about the new service performance, supported by the FLAME benefits and understand their impact. The extraction of measures and the generation of knowledge about the performance of the new services are key outcomes of the project. These service metrics and their link with QoE and QoS parameters are especially profitable for media service providers. Media service providers, which are responsible for the design and deployment of new services, find crucial capabilities in FLAME: a platform for the orchestration of media services; advanced foundation media services (FMS) to re-use; real-life city infrastructures for testing and validation; and suitable measurement and knowledge about their service performance.

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

3 PARTICIPATORY MEDIA FOR INTERACTIVE RADIO COMMUNITIES

3.1 MOTIVATION

VRT Innovation has developed two qualitative user experiments to explore the innovative technical capabilities and the value of the FLAME platform for developing new media interactions inside the city. In the first, we focussed on translating the local media production experience to developers and service providers in order to create seamless quality of service. In the second, we further iterated upon the learnings of the first user experiment and focussed more on the actual storytelling process when producing in the physical world.

In a third "Urban Hacking in 5G" hackathon activity, VRT presented the design and outcomes of these experiments as examples of how a public broadcasting company searches for new interactive and participative media experiences.

As an overall motivation for its work in the FLAME project, VRT explores how it can improve the engagement of city communities in producing media together.

In relation to the media service trends identified in Section 2.1, we specifically aligned our experiments with the desire to provide localised contents and experiences (trend 3) in a context of content consumers becoming content producers (trend 4), while providing an enriched experience beyond today’s media technologies (trend 2).

3.2 VALIDATION EXPERIMENT

3.2.1 Description of Experiment

In the first user experiment, on Mobile Journalism, 12 participants were asked to join a role-play of a fictive local media company. The play aimed to evoke a context where participants experience the challenge of capturing compelling news stories and where they can experience newsroom time pressure. Four teams were asked to create two short news bulletins storifying capturing citizens thoughts on the current and the future City. All participants got an introductive course on mobile journalism and storytelling. The newsroom editors guided the remote journalists using a live chat channel, provided via a mobile application developed by VRT. Captured videos were sent to the newsroom using the FLAME platform.

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

Figure 1. The journalist as a designer. Creating tangible conversation openers used in interviews (Photo FLAME , 2018).

Figure 2. Local news room activities. Incoming interview (Photo Marc Godon, VRT Innovation, 2018).

In the second experiment, VRT deployed an interactive media application in the Millennium Square, situated in Bristol, stimulating a debate on future city challenges. The FLAME platform was tested regarding its capabilities of media distribution involving media encoding, storage, and load balancing at the edge of the network. 5 participants were asked to respond with their own made videos in 2 groups on several sustainability challenges posted at a particular spot in the Millennium Square. Access to media was limited to the immediate surroundings of the square, suggesting a physical experienceable media interaction zone. After this outdoor experiment, participants were asked to envision their own desired media interaction system in a workshop.

Figure 3. Students explore the interactive media trail with Bristol's Sustainable Future plans and UNSDG as themes (Photo Marc Godon, VRT Innovation, 2019).

Figure 4. Participants explain their vision on the city's future and the role of new interactive media to Smart City experts (Photo Marc Godon, VRT Innovation, 2019).

The square was divided into 6 virtual zones, each corresponding to a certain topic. The topics stand for the big issues each transactional city is confronted with. In the edge was related content to each topic stored which could be retrieved using the VRT app, so the participants could get informed. Also, via the app, participants could upload their own contribution to the topic.

The content was delivered via the storage FMS and the adaptive streaming FMS. There was a failsafe so the app could switch remotely between streaming over HLS or using progressive download due to issues with the streaming server during the tests leading up to the experiment.

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

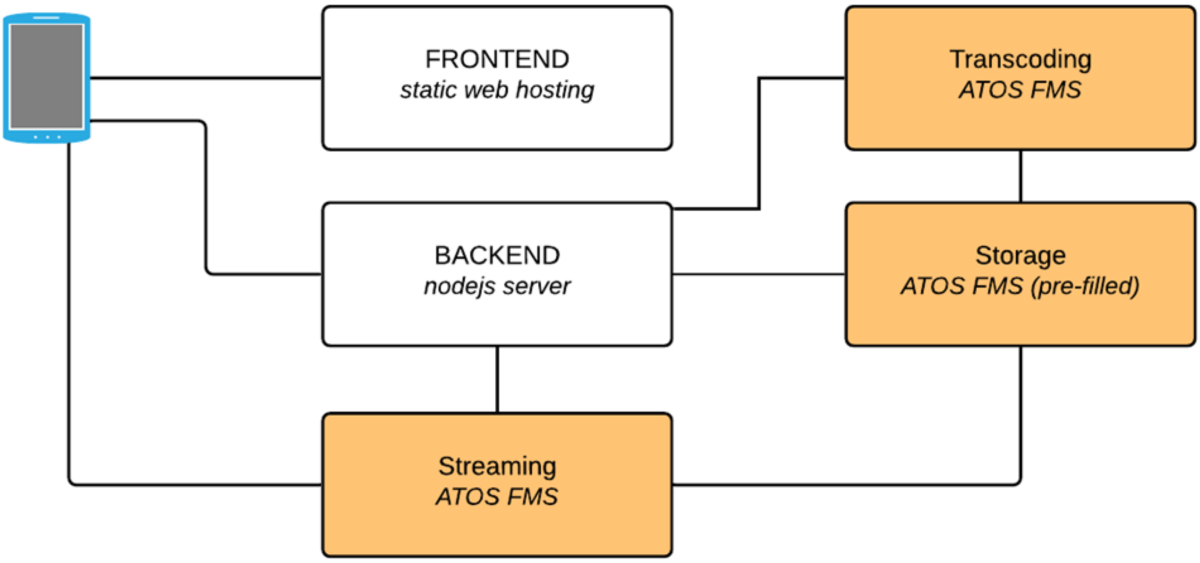

Figure 5. Service function chain used during the second experiment. The backend functioned as a controller service for the web application on the client. Uploaded videos were transcoded in chunks to be served for the adaptive streaming service.

Figure 6. Sign 'Environment' attached to a bench on Millennium Square (Photo Marc Godon, VRT Innovation, 2019).

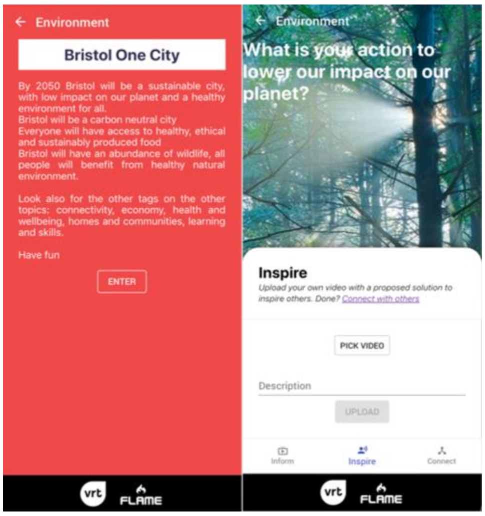

Figure 7. Screenshots of the VR app. We see the corresponding topic 'Environment' (Screenshot Klaas Baert, VRT Innovation, 2019).

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

3.2.2 Results

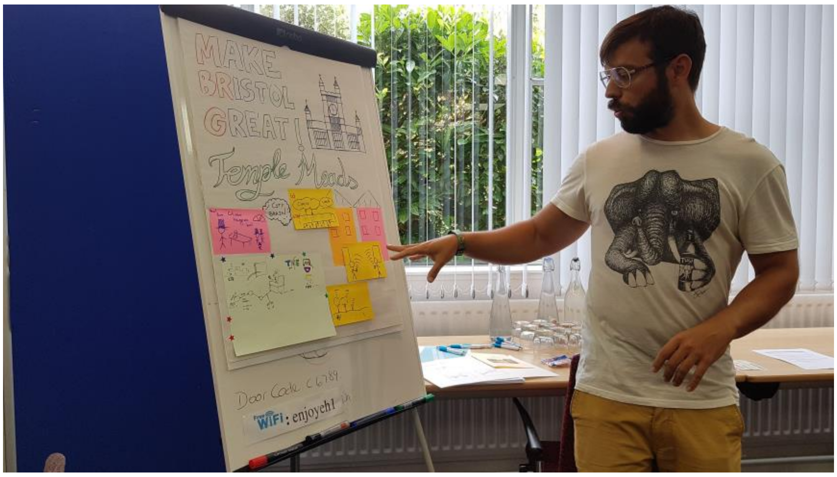

Figure 8. A participant presents his learnings in the first experiment (Photo Sandy Claes, VRT Innovation, 2019).

After the first experiment, the four teams presented their learnings to the group. Two of the three studios asked permission to use a 360-degree camera. In both teams, one person was experienced with this technology, i.e. one vlogger and one entrepreneur in 3D-supported web browsers. We asked them to think about if and how this would fit in their run-down as we were aware that presenting a journalistic item in 360-degree video is not easy, even for professionals. Yet we were open for them to bring in their expertise. All teams mentioned they were surprised how unpredictable the situation in the city was, and how the content of the interviews steered their work, ensuring the need to rework their run-down on the spot. The studios that worked in 360-degree video, both mentioned how they became aware how this technology mostly allows for focus on the surrounding context. Studio 2 reflected on how 360-degree video, instead of being a presentation format, could be a producing tool to help the edit producer to have an overview of the team on site, and spot interesting storytelling angles or additional footage.

Three studio teams mentioned how media technology should support journalists to focus on bringing stories of citizens, and not to be distracted by the technology. Team 2 presented a solution to integrate different tools. For instance, the chat application might be integrated with the run-down of the interview to better allow for sudden changes or new storytelling angles. Also, in this integrated tool, according to team 1, there should also be attention to the ethical aspects of capturing video. Moreover, when filming in 360 degrees, passing citizens are not aware their presence is included, triggering new ethical questions and consents.

In the second experiment, the two teams reflected on their experience through a co-design workshop in which they iterated upon the improvement of the available service.

Some technical issues were experienced by the participants which impacted the QoE significantly. Participants were informed about the hardware and software architecture of the experiment and accepted the explanation regarding the challenging nature of these kinds of technical experiments. Note they are all IT students except of for one mathematics student.

During the workshop, the students were asked to evaluate the experience in group using the ‘I like, I wish, I wonder’ method. Next, they needed to design their ideal and improved application and pitch this idea to the VRT team, assisted by Stephen Hilton, an international recognised smart cities expert.

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

The review from the city expert was very positive:

- Millennium Square as living university or the City becomes the university. This should be pitched to the head of University of Bristol.
- "Bump into media", together, locative, interaction, personalization and adaptation are key mechanisms for going from spaces to places.

The experiment confirms:

- Virtual layer can be an overlay to promote social interaction.
- A physical square can be a virtual forum as long as there is an association or mapping between them.

3.3 INSIGHTS

The first experiment, on Mobile Journalism, revealed the potential value of media creation and collaboration tools, those tools proactively residing in the FLAME edge network, to assist the light equipped journalist in his fast-responsive creative work. One team did experiments with 360-degree media and virtual reality technologies to illustrate the potential value of, e.g., remote directing. Participants agreed on the importance and excitement of collaborative storytelling in the City and the positive influence of high quality of service networks on distributed and collaborative media production workflows enabled by the FLAME platform. The first experiment confirmed the valuable idea of a virtual editor room, together with smart tools to support the creative act of making media, deployed in the edge of the FLAME network.

The second experiment, on the Millennium Square, demonstrated the value of deploying an interactive media experience in a public place. Using this square as a *mise-en-scène* worked very well. According to a historian: this should not be surprising. Public places, such as the Millennium Square, have always been the ultimate meeting place for commerce or entertainment. They are giving way to meet other citizens or visitors and exchange stories. In the experiment, we link a digital social network with a physical city space – in a meaningful way. It fits, and it adds value. A close integration can be done using the different service components deployed on the FLAME platform. Public broadcasters should explore more this potential symbiosis between local media creation and city life.

One aspect, common to all experiments is the observation of the ease of imagination, the fluent incorporation of design thinking and media creation techniques, and the civic mindset of the proposed applications leading to the sharing of a better understanding of their city[1]. This could confirm the existence of a positive interference between two systems. "The City", continuous in transition with its need for innovations, and the FLAME platform which is supporting the easy creation of local interactive media experiences.

[1] https://www.ict-flame.eu/news/citizen-services-gaming-and-immersive-media-showcased-on-a-new-5g-platform-in-real-life-trials/

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

As stated in D5.2 the experiments of VRT conceptually focus on both sides of the news delivery process, namely *ingest of content* and *broadcast of content*.

What the experiments showed is that at a conceptual level, the ingest of media into the edge has proven to be meaningful. While broadcasting content can make use of the platform's multicasting capabilities, uploading content remains, naturally, a single connection operation. The benefit of using the platform for content ingest lies in the use of future media services.

Because of the timing of both experiments and the readiness level of both the platform and the future media services, there was no technical outcome from our experiment in the latency debate. However, this is being explored in the platform benchmarking tests reported in FLAME D5.6 [11].

It is clear to say that performing content analytics in the edge to filter the uploaded content before transferring to the cloud, is more effective in terms of data transfer. Although we lack a real cost model at this point, we can expect that this will have a positive influence on the cost, as transferring data from A to B is the highest in the whole media chain. Also transcoding content first, thereby compressing and reducing the amount of traffic from that point on is better to be done as early in the media chain as possible.

The experiments made it clear that there is a need from media companies on the ingest side of the media chain and this need will only grow in the future with more media consumers becoming media producers themselves.

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

4 VALIDATION EXPERIMENT: PERSONALIZED MEDIA MOBILITY IN URBAN ENVIRONMENTS

4.1 MOTIVATION

During the last years we have assisted to the huge growth of services that make video pervasive in homes, public spaces and also on the go using personal and portable devices. Big broadcast companies have invested in developing apps and services to make their contents available to their clients in any moment and in any situation (e.g. Sky Go, DISH Anywhere, Netflix, Amazon Prime Video, NowTV, etc.).

In the meanwhile, users experience a lack of solutions for accessing personal videos (e.g. digital content in home-based Video on Demand (VoD) platforms or video-surveillance recording systems) ay *anytime anywhere they want*. The realization of such services is dependent on access to the personal home network. As a direct consequence of that, also concurrent fulfilment of private contents by multiple users is limited by the personal home hardware and network performances.

Addressing those needs, the Personalised Media Mobility (PMM) scenario experimented on top of the FLAME infrastructure in the city of Barcelona. It stressed the attention to Personalisation, Interaction, Mobility and Localisation (PIML) aspects of the media distribution in a Smart City.

The PMM experiment evaluated how the FLAME platform allows media service providers to serve users their personal content while they are moving about within the topography covered by FLAME network. In particular, the developed PMM scenario exploits these key features offered by the FLAME platform:

- CLMC alert triggering
- Intelligent service endpoint management
- Automatic load balancing and consequent QoE improvement
- Dynamic service routing to direct traffic to the most appropriate local service instance
- Reduction of network traffic through the localization of traffic flows source and destination, wherever possible, also addressing the aforementioned latency reduction

In relation to the media service trends identified in Section 2.1, we specifically aligned our experiments with

- *Trend #1. Demand for improved quality,* including more pixels (in our case up to 4K resolution for streaming at home) at different resolution and encoding quality;
- *Trend #5. Users wish to watch any content anywhere and at any time*, e.g., as they navigate across a city or during a daily commute, with particular focus on personal media contents.

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

4.2 VALIDATION EXPERIMENT

4.2.1 Description of Experiment

The PMM offers an experience in which “my content is always available” from anywhere (within the smart city), from any device and for anyone with access to personal media server (typically family members).

In our experiment we presented two scenarios:

- **PMM distribution in walking areas in Barcelona** (see Figure 9): the main goal is to validate how users can be able to start, stop and resume the playback of audio/video content from another place with another device. The consumption of the media content continues while the user is walking in the Smart City without losing bookmarks and preferences;

- **PMM for multiple concurrent streams within the Smart City** (see Figure 10): the QoE offered by the platform is tested with multiple users, dislocated along the smart city area, accessing to the same personal contents.

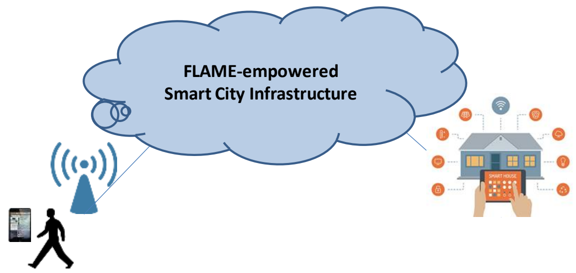

Figure 9: PMM Scenario 1 -Distribution of personal media in walking areas in Barcelona

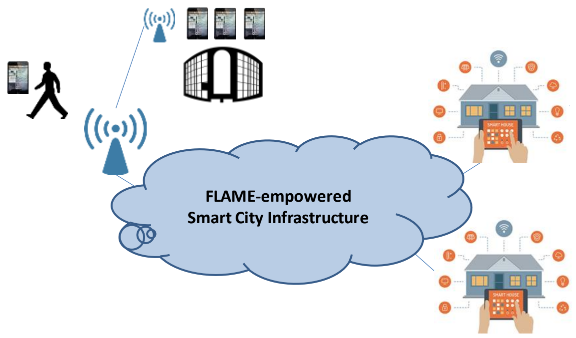

Figure 10: PMM Scenario 2 - PMM for multiple concurrent streams within the Smart City

The PMM experiments over the FLAME platform evaluate functions for media distribution that are suitable for serving dispersed endpoints in the FLAME-empowered Smart City. Indeed, the evaluation is focused on how FLAME allows to go beyond the traditional Content Delivery Network (CDN) architectures currently available for media distribution over IP.

As shown in Figure 11, the PMM Service Function Chain (SFC), when the service is at first deployed, consists of a Personalised Media Origin Server (based on PLEX) deployed in the core data centre of the FLAME infrastructure in Barcelona. Initially, the Origin Server is active and connected through the FLIPS routing capability in order to be reachable by the Access Points of the FLAME platform in the area of Pere IV district. Replicas of the Origin server are placed but not connected into “gateworks” clusters (street cabinet units for edge computing).

The PLEX Origin Server periodically sends data about streaming statistics (bandwidth, number of active streaming, type of active streaming, etc.) to FLAME’s CLMC.

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

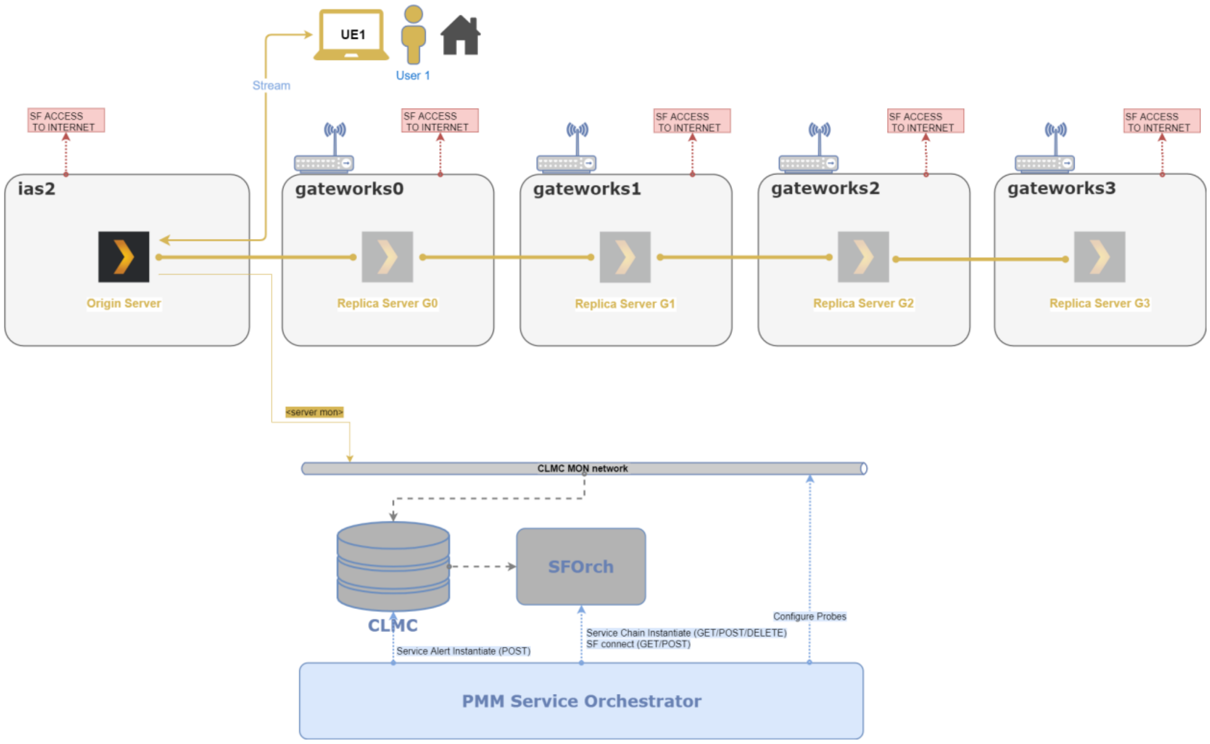

Figure 11: PMM Trial Barcelona – Initial Deployment

Once three or more concurrent streaming sessions are detected on the server, the CLMC triggers the scaling-out of the PMM SFC (see Figure 12).

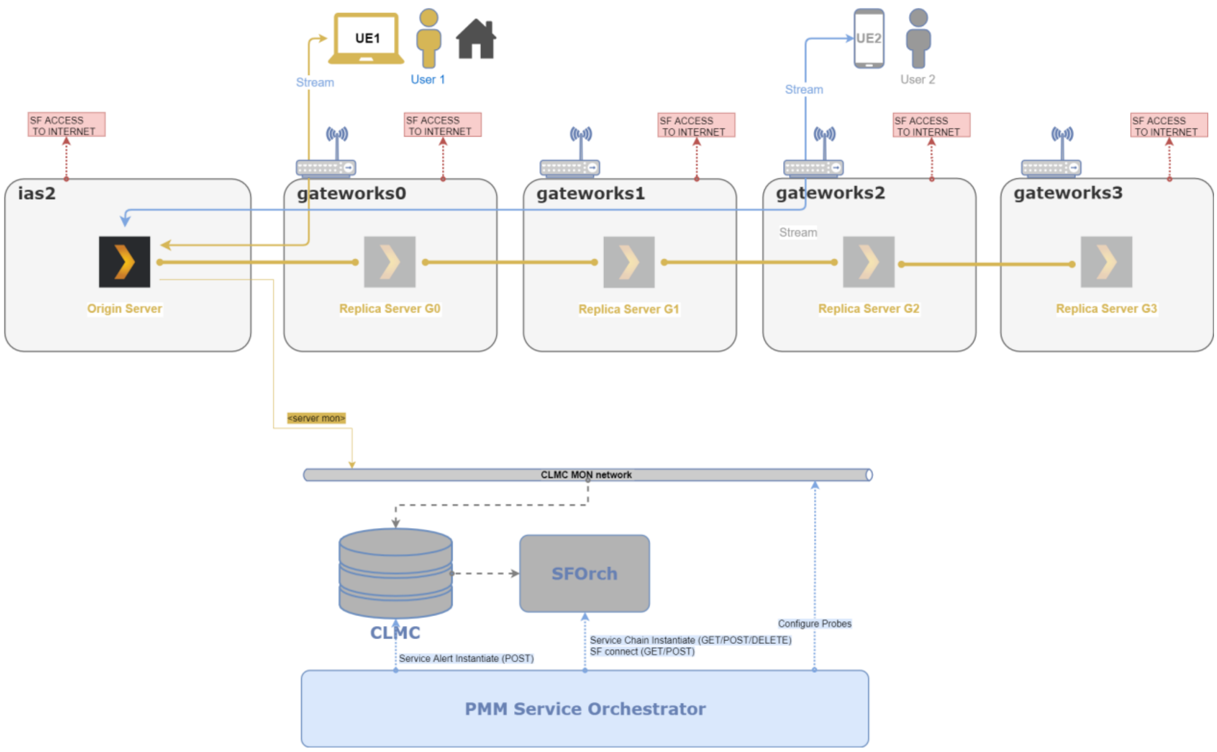

Figure 12:PMM Trial Barcelona – Scale-out trigger condition

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

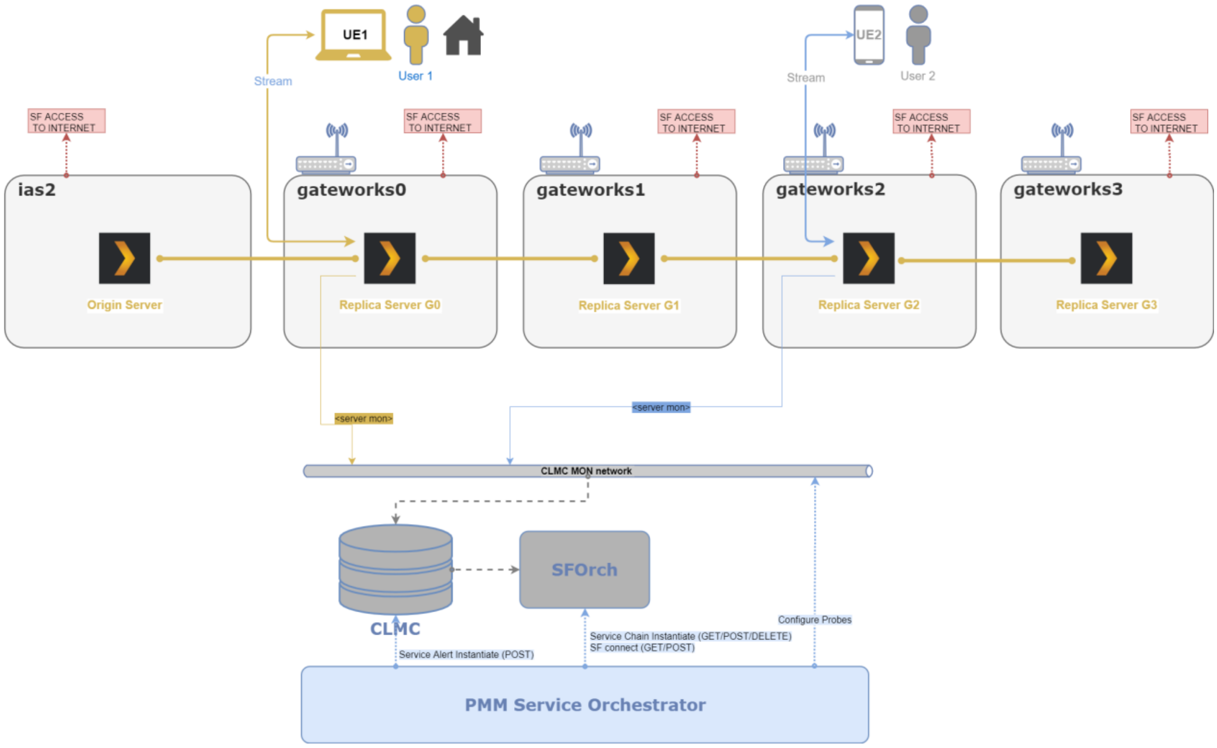

Figure 13: PMM Trial Barcelona – Scale-out configuration

After the scale-out has been processed (see Figure 13), the PLEX Origin Server and its replicas result in "active" mode on all the four clusters and the core datacentre. Each Origin Server instance is then reachable via the Service Function Routing (SFR) component [10], which determines which is the most convenient route towards the activated replicas in the various edge points. Consequently, each user is automatically served by the closest server and the load balancing among the Plex Origin Servers is managed automatically and dynamically by the FLAME platform.

In the meanwhile, while the service is running and media streams are activated by the users, monitoring data is collected from the Origin Servers in order to be processed by CLMC. The analysis of the collected data eventually can trigger the scale-in alert; specifically, this operation is automatically performed when no connections are detected, restoring the initial state of the deployment, when only the PLEX Origin Server in the core datacentre is "active".

The PMM Trial execution has been articulated three Phases:

1. Dry Run without FLAME Scale-out function activated to verify correct operations and profile "no-scale out" QoE
   (Conducted by i2CAT and NXW personnel)

2. Dry Run with FLAME Scale-out function activated to verify correct operations and profile "scale out" QoE
   (Conducted by i2CAT and NXW personnel)

3. *Public trial to make a stress test of the platform and the PMM service and to collect data for QoE post processing and analysis*
   (12 users and pre scaled-out service)

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

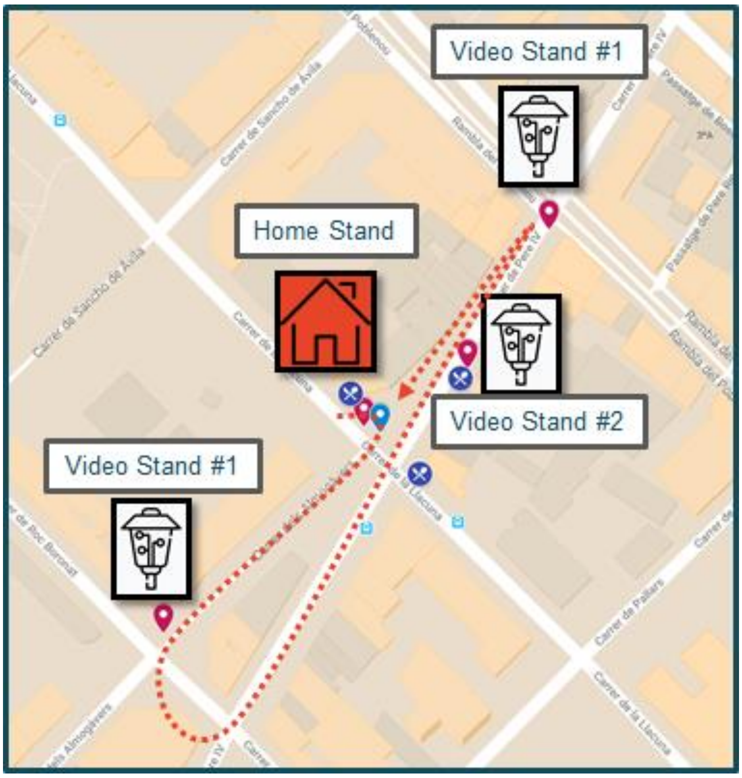

Figure 14. PMM Trial Barcelona - Map and checkpoints.

As shown in Figure 14, the emulated “Home” location was placed in Bar Llacuna (Carrer de Pere IV and Carrer Llacuna corner). At the Llacuna bar, users were asked to register to the trial, to compile the consent form and finally to check their UE terminals. Then instructions and credentials for accessing the Service were given and users were guided through the installation of the screen recording utility[2] and the Mozilla Firefox browser for using the Plex web streaming client.

For running the trial, three groups of 4 people were created. Users of each group were asked to first start a video recording, show their current position opening Google Maps and then to start streaming the first selected video. The first stage allowed experimenters to check the correct configuration of each device, then users (always divided in groups) sequentially visited the three checkpoints (Video Stand #1, #2 and #3) and, in each of them, they streamed and recorded a planned video.

After visiting all the video stands, they came back to the “Home” site to share with the organizer their recordings. The Nextworks team then took care of the post processing of the data by analysing CLMC statistics and by performing an estimation of the QoE[i] of the collected videos with the VQ tool[3].

On the Origin Server each streamed video was made available in three versions (4k, 1080p and 720p) to meet most people’s device requirements without the need of transcoding.

---

[2] ScreenCam for Android, OBS for Windows

[3] Video Quality (VQ) – See references [3] [4] [5] [6] [7] [8]

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

4.2.2 Results

The first step of the trial consisted in a dry run conducted by NXW and I2CAT personnel. Users started investigating the SF working without using the optimization instruments offered by FLAME platform: the service was deployed leaving connected only the Plex Origin Server placed in the core datacentre, which served all the incoming streaming requests.

With this configuration, it was not possible to stream more than 2 videos concurrently. Starting the third one, data traffic immediately stopped, making impossible to continue the service fulfilment for all the 3 users. It was also evident that a single transcoding session (due to the client requesting the video in a format not available on the server) pushed the CPU activity of the server over the 80%. Concurrent transcoding executed on a single server is therefore impossible without degrading the streaming experience.

From top-left, Figure 15 shown in sequence:

- Active streams (line 1)
- Overall data traffic (line 1)
- Requested bandwidth on the Origin Server (line 2)
- CPU usage (line 2)
- Data traffic per datacentre, it overlaps with the overall data traffic, since just one Origin Server is connected (line2)

Figure 15: Dry Run phase 1 (no scale out function)

After having collected these insights the Alert descriptor was refined and uploaded, enabling the SFC able to take advantage of FLAME platform functionalities. It was then possible to validate the correct behaving of Scale-out and Scale-in operations and to give a first evaluation of the advantages derived by the scaled-out configuration.

Figure 16 shows the platform behaviour while performing the Scale-out after a third stream was started. Then, with all the Plex Origin Servers connected (one on each cluster), users were able to stream up to 6 videos concurrently. At the end of the streaming sessions, when no user was any longer streaming for 20 seconds, the Scale-in was performed bringing the SFC back to its initial configuration.

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

Figure 16: Dry Run phase 2 (scale out function activated)

After the Dry Run stage certified the correct operations of the service, a pre scaled-out SFC was deployed in order to run the public trial and evaluate the performances of the service over Barcelona infrastructure.

Figure 17: Public Trial Active streams over gateworks

In Figure 17, it is possible to observe the statistics collected from the different groups while streaming the media contents at the different Video Stands. The pictures show how the various streams have

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

been activated on the different available replica serves (total number of users from replica servers in Figure 17 top left corner), and how they distributed across the various replica serves emulating home (no user in iaas2 server in Figure 17 top right corner, users distributed in the various gateworks 0-3). Graphs show in blue lines the direct streaming play and in orange the activation of transcoding in the gateworks servers, with subsequent degradation of service quality.

Here below are reported the highlights registered running the trial:

- Group # [Windows 10 tablets]:
  - FLAME2 (GATEWORKS2) - (home) 4/4 were able to fully load the video
  - FLAME3 (GATEWORKS3) - (stand #1) 1/4 were able to fully load the video
  - FLAME1 (GATEWORKS1) - (stand #2) 4/4 were able to fully load the video
  - FLAME0 (GATEWORKS0) - (stand #3) 4/4 were able to fully load the video

  Note: one user was not able to record all the videos due to exhausted disk space on his device

- Group #2 [Android devices (4 smartphones)]:
  - FLAME2 (GATEWORKS2) - (home) 2/4 were able to fully load the video
  - FLAME3 (GATEWORKS3) - (stand #1) 1/4 were able to fully load the video
  - FLAME1 (GATEWORKS1) - (stand #2) 0/4 was able to fully load the video (2/3 of video)
  - FLAME0 (GATEWORKS0) - (stand #3) 0/4 was able to fully load the video (1/2 or less of video)
- Group # [Android devices (3 smartphones, 1 tablet)]:
  - FLAME2 (GATEWORKS2) - (home) 3/4 were able to fully load the video
  - FLAME3 (GATEWORKS3) - (stand #1) 2/4 were able to fully load the video
  - FLAME1 (GATEWORKS1) - (stand #2) 2/4 were able to fully load the video
  - FLAME0 (GATEWORKS0) - (stand #3) 3/4 was able to fully load the video

  Note: one user could not record 2/4 videos due to exhausted disk space on his device

Observations:

- In certain moments, not all the data packets from the streaming sessions were delivered. Group 2 experienced it in FLAME1 (GATEWORKS1) and FLAME0 (GATEWORKS0) where the reproduction of the content stopped and never resumed. Users tried twice - cleaning Firefox's cache and reconnecting to Plex Server – but the same happened again and roughly the same amount of video segments were buffered before the video streaming stopped. Group 3 reported videos getting stuck at the end (with full buffer), the video was entirely played but did not reach the end with few video segments missing. The cause for the experienced block in filling the application buffers is not fully clear at the moment. We think it could be related

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

to either temporary WiFi bottlenecks due to the number of concurrent traffic flows on the FLAME access point and/or to a potential deadlock among the active replica servers occurring during the synchronization of streamed video chunks towards the same terminal.

- User experience was heavily dependent on the user device, in particular group 3 had an old Android device with slow CPU, small amount of RAM (2GB), and less than 5" screen size. With small screens, the transcoding function on the server was activated instead of performing a direct play. This affected other concurrent streaming due to the server overload.

- Having the screen recording activity running in background reasonably did not affect the quality of the streaming, since all registered issues were connected to data packet availability and not to responsiveness and decoding time.

- The infrastructure provided availability of three cores and 4 GB of RAM on each gateworks. This amount of resources was not enough for media applications, in particular when transcoding was required.

- The SFC doesn't perform multicasting since the content is one single hop away from the end user and each one might require a different quality version of the content depending on the device capability.

- Each user was served by the Plex Server installed on the cluster connected to access point the user was attached to, relying on the output of ''shortest path'' algorithm. This way, the user was served by the nearest Origin Server replica, reducing latency, but no load-balancing was performed among all the available servers.

At the end of the trial, users shared their video recordings with the experimenters that took care of the QoE analysis of the captured videos. The measurement of the quality of experience has been performed by using the VQ software and calculating the value of 5 relevant video metrics[4] that are described in the following table.

Spatial activity This metric resembles the amount of spatial details in the image. Too strong compression would result in diminished value of this parameter.	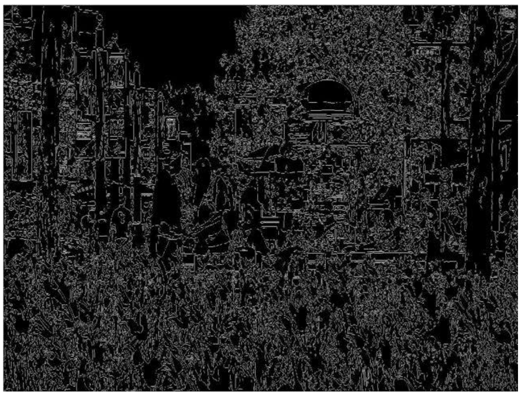	No distortion from 0 to 60

[4] 5GINFIRE WIKI: http://wiki.5ginfire.eu/experiments/RobotView5G

Temporal activity This metric corresponds with an amount of temporal activity in the video. The higher the values, the bigger the change from frame to frame.		No distortion from 0 to 20
Blockiness This value detects distortions characteristic to video compression algorithms operating on (macro)blocks. When such an algorithm is allocated too little bitrate, the produced video characterizes with blocking artefacts this metric is intended to detect. The lower the value, the stronger the effect.		No distortion from 0.9 to 1.01
Blockloss This parameter indicates the amount of outdated content present in the picture in the form of distinct blocks. These can appear when data loss is present.		No distortion from 0 to 5
Noise This metric detects noise, randomly modifying single pixels. Such noise is characteristic to digital video cameras working in dark conditions.		No distortion from 0 to 3.5

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

The graph in Figure 18 shows in green the amount of the frames with no degradation over the total number of frames analysed. Most of the corrupted frames (71%, see Figure 19 and Figure 20) presented degradation of a single parameter (in particular Blockness). A detailed plot of the distribution of the amount of frames affected by blockness is reported in Figure 21.

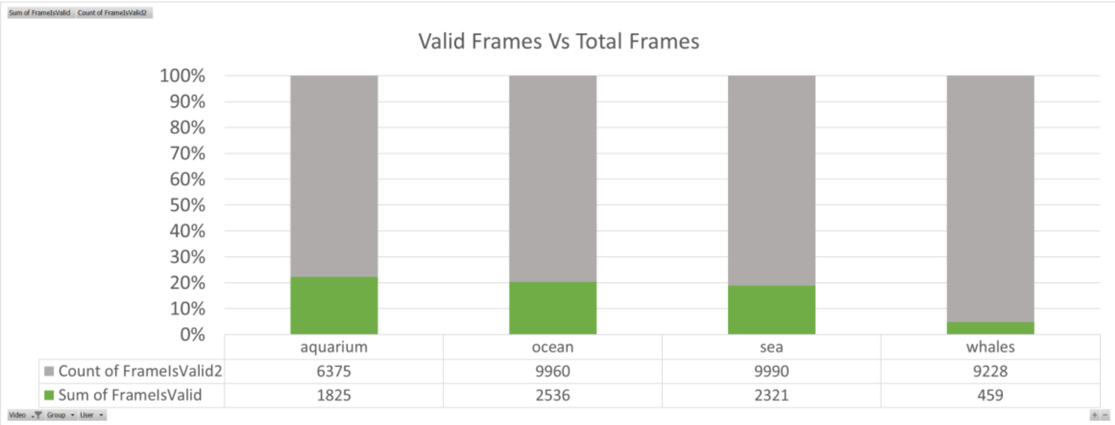

Figure 18: PMM Trial Barcelona – Valid frames over total number

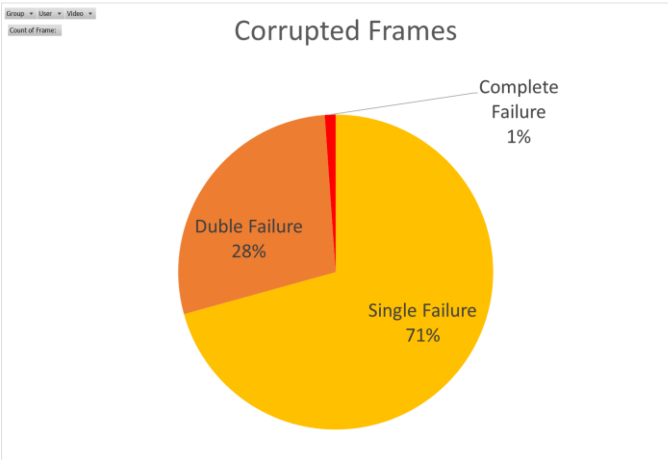

Figure 19: Percentages of corrupted frames

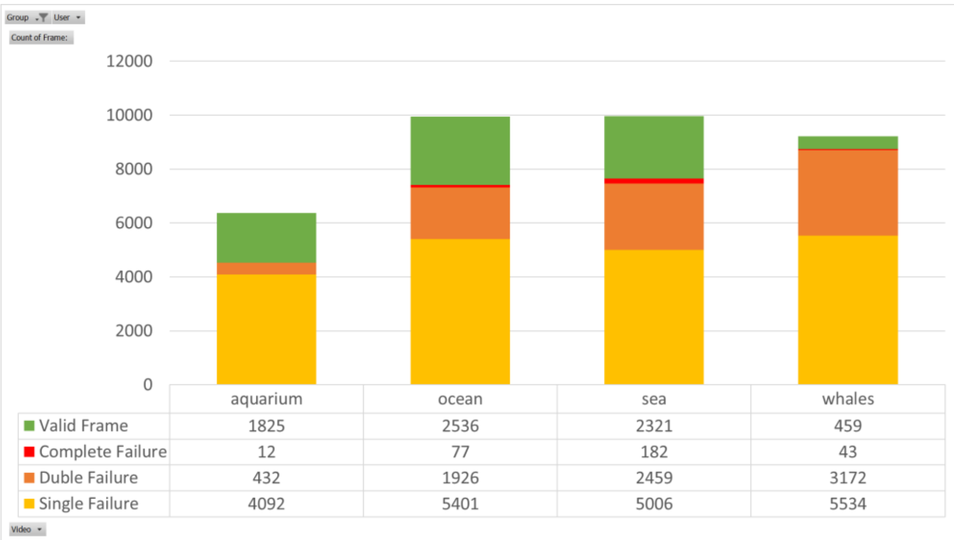

Co-funded by the Horizon 2020
Framework Programme of the European Union

Figure 20: Frames divided per distortion degrees

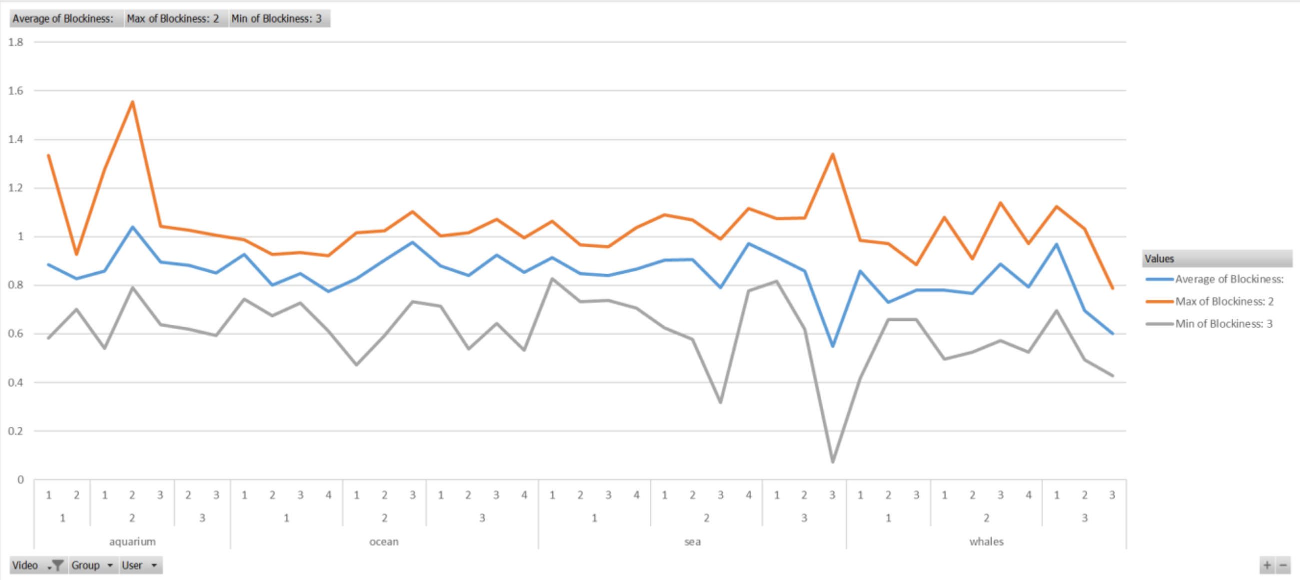

Figure 21: Blockness frames, max, average and min values

4.3 INSIGHTS

As initially planned in deliverable D5.2, this PMM experiment has been designed to identify relevant technical feasibility aspects and business implications of distribution of media in smart city areas. The actual experimentation occurred in Barcelona has allowed us to derive useful insights for the main business actors of interest, i.e. Media/VoD service provider and Media/VoD technology provider but it allowed also to derive useful insights for the infrastructure owner/operator, as detailed in the following.

From a technical perspective, the PMM experiment has emulated the behaviour of a group of people sharing a media streaming service where personal contents are stored and willing to consume these media while moving from a central place (home) into a specific area of the city. To emulate the localisation of the group of people in FLAME coverage, we configured our scale-out function to activate and connect all the placed Origin Server replicas in the three remaining edge cabinets. The implementation in a wider geographical area, e.g. in the city of Barcelona, would require using localisation functions to activate only the server replicas available in the district where users are identified and may require coverage along their walking paths in the area.

The dimensioning of the experiment documented in the previous sections shows that the deployed media SF endpoint could not support the coexistence of 3 groups of 4-5 people each, with a high number of blocked streams, due to transcoding on the target media server (origin or replica). The analysis of results shows that the main cause for this behaviour can be the limited amount of computations resources made available for PMM SFEs in the Barcelona platform (2 virtual CPUs and 4GB RAM per media server). In fact, PMM media servers are generally dimensioned with 4-6 CPUs and 8-12 RAM to serve 4-10 users. Moreover, the use of different devices with different form factors (tablets, smartphones with and without FullHD, various form factors for the screens) caused a number of parallel transcoding activated by the media application to adapt the played content to the actual device capabilities. Despite the creation of pre-transcoded versions of the same contents which was

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

aimed to avoid the transcoding event and make use of direct streaming, we experimented recurrent degradation of the overall QoE both in terms of objective parameters and if subjective user "acceptance" of the media service quality when run on FLAME.

Despite the performance issues due to limited hardware resources made available to serve >10 users, it is worthwhile to note some experimented benefits for the Media Service Provider using FLAME. In fact, the PMM designers experimented with how the FLAME platform can really offer easy and intuitive mechanisms for designing and deploying new services from a Vertical industry perspective, i.e. with very limited knowledge of the platform and network mechanisms and technical solutions (SDN, SFC, service stitching, flow routing, etc.). ***The lower complexity of the FLAME service descriptors compared to similar descriptors needed to implement similar functions in NFV orchestration platforms based on ETSI MANO specifications*** clearly shows significant levels of beneficial abstractions for the Vertical who simply focuses on declaring service elements and connection status, thus skipping a non-trivial considerable amount of configuration required in NFV platforms to specific how these functions have to be stitched in a forwarding graph, how the network functions are profiled, etc. Moreover, the mechanism by which FLAME allows to define policies and alerts is intuitive and can be easily adapted and tailored depending on the available computational resources in the target environment as well as the design of the service chain can be extended and enhanced for covering wider Smart City areas.

From a business perspective, it has emerged clearly that a considerable investment in infrastructure resources has to be made to prospect a feasible PMM service in Smart Cities. The trial experiment demonstrated that a media server SFE dimensioned with 2 CPUs and 4GB RAM. This data can be used together with the target number of parallel SFCs to be supported (I.e. families to support in a given area/district) to dimension the physical computing and network infrastructure to adopt.

To summarize insights and put them in in relation to the expected outcomes from this experiment documented in deliverable D5.2

Table 1 – Summary insights on PMM experiment in Barcelona

Expected outcome from PMM experiment as planned in D5.2	Results
[Technical] A better understanding of the mechanisms to optimize bandwidth and resources among the core and edge parts of the city infrastructure	Achieved. FLAME platform configuration allowed easy mechanisms to test service deployments and scale-out/scale-in policies via simple abstract service descriptors. This was further enabled through the provided devOps toolchain that FLAME provided in the form of sandbucket, sandpit and real-life test bed.
[Technical] The assessment of the feasibility of personalized media streaming services in the software defined infrastructure.	Partly achieved. FLAME infrastructure in Barcelona allowed to find some actual dimensioning data for media servers to be adopted in for a PMM service. However, the degradation of service performances and the limited number of concurrent streams supported (due to limitation in available computing resources in the trial) cannot allow to complete an assessment of feasibility for such a service

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

[Technical] An evaluation of the benefits of the FLAME benefit through the capabilities of the FLAME SFR component, i.e., flexible routing among surrogate functions with respect to traditional Content Distribution networks	Partly achieved. FLIPS worked as per specification and automatically routed traffic towards the available surrogate functions which were activated after scale-out. However, in case of service degradation the limitation on resources impacted user's perception on the quality of the service deployed. The experimented blockage in fill of client application buffers might be related to deadlock conditions generated by the autonomous selection by FLIPS of the target surrogate function, with a consequent mismatch (at server side) among media chunks just sent to client and to be sent.
[Business] The user acceptance and interest in such PMM service, through the trial phases in public areas.	Failed. QoE degradation conditions discussed above generated unsatisfaction in the users towards the usability of the service they experimented.
[Business] The applicability of a PMM service model to Smart Cities, e.g. to stimulate new offers for value-added touristic services by the municipalities or for enhanced user information streamed via various devices and access points (e.g. digital signage totems)	Partly achieved. The size of the experiment run in Barcelona does not allow to derive conclusive information on this aspect.

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

5 VALIDATION EXPERIMENT: COLLABORATIVE INTERACTIVE TRANSMEDIA NARRATIVES

5.1 MOTIVATION

This validation scenario demonstrates how the FLAME platform supports the delivery of location-based augmented reality (AR) stories within an urban environment. The scenario augments physical locations allowing mobile users to experience interactive and content rich spaces that bring together digital and physical worlds. This requires new and advanced media services that are localized within the physical environment. The FLAME platform enables intelligent orchestration and management of these media services, for example to support localized media processing or delivery capabilities.

The location-based AR storytelling (AKA. City-wide Storytelling) system presents multimedia content as part of the storytelling process. Some of the story content is video, audio or text, and a special emphasis is placed on delivering virtual 3D renderings of animated characters and props in mixed reality environments. A primary design concept is to not require the mobile application to have all content to be preloaded. Instead, a lightweight mobile application downloads content that is contextually appropriate based on both the physical location and the state of narrative progression of the user.

The aim of this validation experiment is to tell interactive stories through AR within a city environment. The user enters the city, opens the FLAME-enabled application, and chooses from a variety of stories that can be experienced in this region. Such stories can be used for tourism, education or simply for entertainment. By augmenting the world, experienced through the smartphone camera, with 3D animated objects, a new level of immersion is reached and could, for example, bring historic scenes back to life in front of the user's eyes.

5.2 VALIDATION EXPERIMENT

5.2.1 Description of Experiment

The validation scenario experiments take place on Millennium Square in Bristol, UK, where a few famous Bristolians are present in the form of life-sized bronze statues. The story is acted out on different AR stages. These AR 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 visualized as soon as the corresponding AR marker is found. The AR markers are used to properly orient the stage relative to the real environment, which brings together the digital and physical world.

5.2.2 Results

Note: at the time of writing the results were still being processed.

5.3 INSIGHTS

Note: these are preliminary insights based on initial results. This section will be expanded upon when the full results are available.

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

The FLAME network allows to spatially distribute media assets within a city, such that each asset is available near the location where it will be consumed. For a media service provider, this enables localized control of latency and bandwidth to reduce retrieval time. For the media consumer, the media content does not have to be pre-downloaded on to the device but can be streamed gradually as the story unfolds.

The FLAME platform offers intelligent resource management capabilities. The scalable nature of the FLAME platform can cope with high user activity, by automatically starting new media content services throughout the city to distribute the load over more instances when latency is reaching critical levels at certain regions. The FLAME platform enables the location-based AR storytelling system to minimize platform utilization (e.g. by turning off or removing a media service) based on predicted local service requirements. The intelligent lifecycle management of localized media services results in more efficient resource utilization.

The FLAME platform could be utilized to integrate additional low latency media services to enhance the immersion and interaction experienced with location-based AR stories. For example, high quality 3D renderings could be generated by powerful servers at edge locations and delivered directly to the client mobile device. Another example is to use low latency state synchronization of several devices connecting to the same edge location. This would enable a multiuser shared experience of an interactive story.

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

6 VALIDATION EXPERIMENT: AUGMENTED REALITY LOCATION-BASED GAMING

6.1 MOTIVATION

Urban environments are not only places for working, commuting, and socializing. They can be repurposed for playful behaviour: for instance, squares become stages for flash-mobs, obstacles become challenges for Parkour runners, and lines on the pavement become traps we should not step on. This kind of city-gamification has been practiced for a long time and, using modern technology, can be partially shifted into the virtual world. Augmented reality (AR) technologies enable mobile devices to overlay fantastic 3D models on urban areas and enable virtual interactions with them.

Even though city-wide games are still a niche market, with the continuous success of mobile gaming, the widespread adaptation of city-wide gaming is only a question of time. Smartphone games are already becoming increasingly sophisticated, both in terms of gameplay and visual quality. The visual beauty, however, often comes at the cost of increased data storage and processing requirements. For instance, often hundreds of high-resolution textures and 3D models have to be stored and processed on the device.

The storage problem can be mitigated by streaming the required 3D assets from a server to the mobile device and persisting them in volatile memory until they are no longer needed. With the increased capabilities of 5G, such scenarios are no longer a mere vision. An additional advantage of streaming 3D assets on demand is that the look of the game can easily be changed: Depending on the user's location, the time of day, season or other metrics, the server can provide 3D assets specific to that environment.

FLAME offers capabilities that can make this vision a reality. By having several services distributed over a city, the data downloaded to the device can easily be adapted to the local area. A service running in the old town of the city could return more dated-looking 3D game assets than a service running in the new shopping district. To deliver game assets as quickly as possible, only low latency and high bandwidth are required. However, efficiency is another important factor to make the application scale well for many users. In a multiplayer AR game, several players may need to download the same game asset at the same time from the service. This is where the Opportunistic Multicasting capabilities of FLAME can reduce the load on the database server by querying the database only once and then distributing the fetched asset to all simultaneous requesters.

This validation experiment will explore how to stream 3D game assets to several client devices with the FLAME platform and what effects Opportunistic Multicasting has on an AR urban arena multiplayer game.

In relation to the media service trends identified in Section 2.1, we specifically aligned our experiments with the user's desires for more engaging, immersive and exciting experiences. The gamification of urban spaces through AR aligns well with this.

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

6.2 VALIDATION EXPERIMENT

6.2.1 Description of Experiment

The Game Technology Center (GTC) at ETH Zurich, Switzerland, has been developing a team-based augmented reality (AR) urban arena game called Gnome Trader to validate the FLAME platform.

The goal of the game is to collect as many tree seeds as possible for your team. Tree seeds can be gained either by planting trees in your AR garden and harvesting them once they're fully grown, or by going to the virtual AR gnome shops distributed across Millennium Square to buy tree seeds. Each tree has its own associated 3D model, which is downloaded when its seed is planted by a player. A planted tree will trigger the download for every other player. Multiple players requesting to download the same 3D asset from the server will trigger multicasting. Some impressions of the game can be seen in Figure 22.

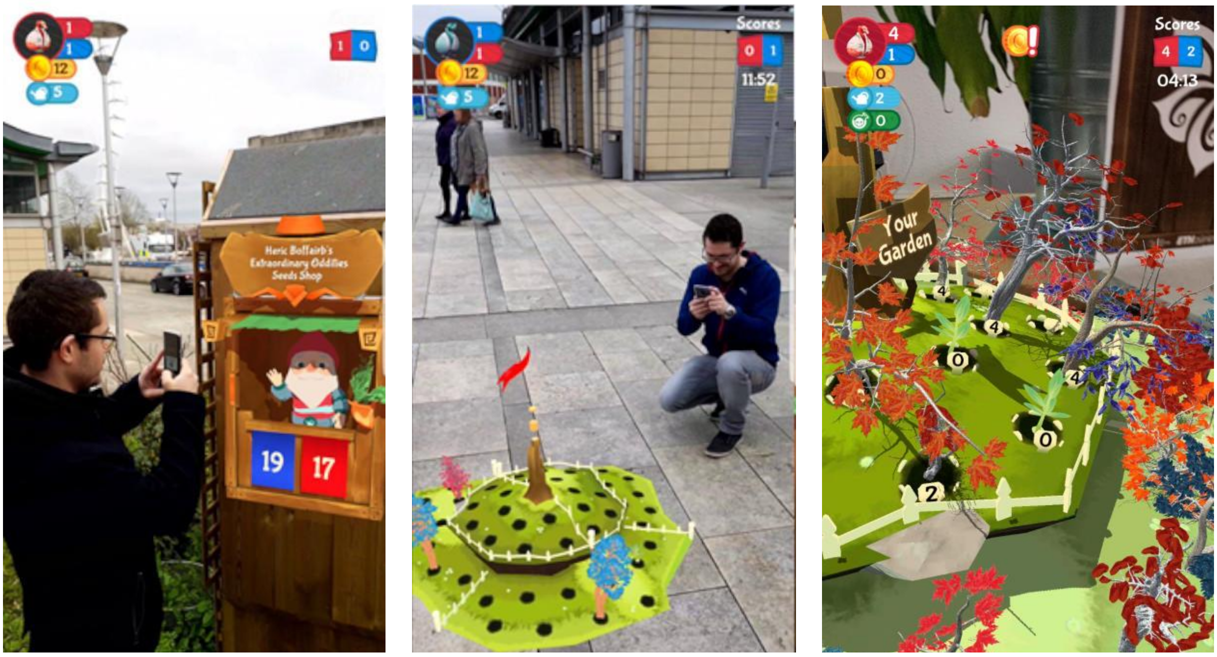

Figure 22: AR gnome shop (left). AR garden (middle). Different tree models in AR garden (right).

Setup

To evaluate the implemented game and game services in FLAME, the GTC conducted several trials with 15 participants on the premises of the Millennium Square in Bristol. The SFC for the experiment contained one storage FMS provided by Atos, which acted as a 3D content database for the game and was assigned 2 CPUs and 4 GB of memory. This database held 100 different 3D tree models, each being roughly 700 KB in size. The game's state synchronization was handled through a Google Firebase database outside the FLAME network to validate the quality of Internet connectivity. We conducted several play-tests with varying number of users to observe differences in QoE and QoS.

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

6.2.2 Results

6.2.2.1 QoE

Since the main goal was to stream 3D assets from a server to the players while not interrupting or slowing down a faced-paced gameplay, the most interesting QoE metric is if the user can notice any delay when a tree asset is supposed to be visualized. If the user cannot notice any delay, then the illusion of having all assets already on the phone is upheld.

Asset Loading

In general, users did not experience any delay when 3D tree models were downloaded from the storage FMS. The only exception were users who joined the game after it had already started. These players had to download all already planted trees at once, which resulted in tree models suddenly appearing in their garden over time. Due to this, the joining of a new player mid-game manifested itself in spikes in the amount of data that is sent from the database to the players. This can be seen in Figure 23, which shows a data transfer spike around 8 minutes and 38 seconds of gameplay when a new player joined.

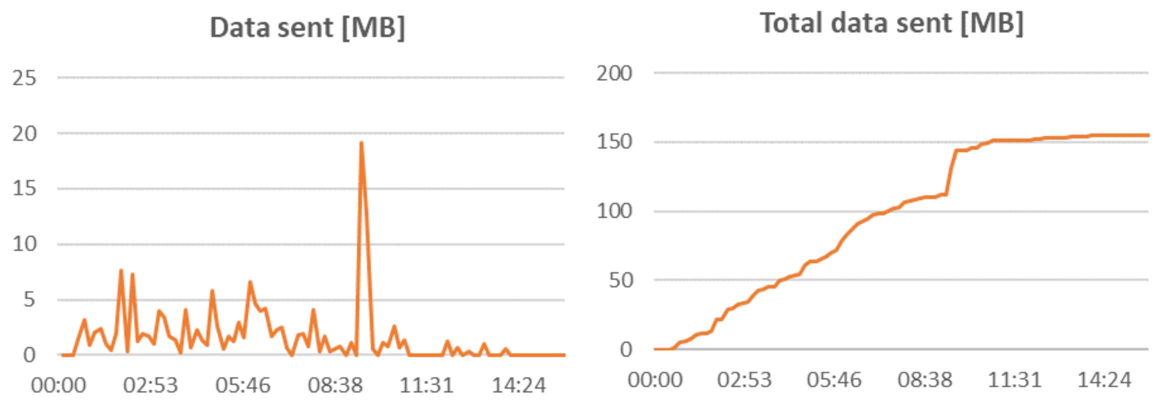

Figure 23: The spike shortly after the 8:38 minutes mark is due to a new player joining.

The size of these spikes highly depended on the number of already planted trees at the time the player joined. Hence, on average, these spikes were larger for games with higher number of players because such scenarios allow for more trees to be planted at any given point in time.

The FLAME platform handled asset loading successfully. Even data transfer spikes were not noticed by players and the game always ran smoothly. Only players who join the game late might experience some delay in asset loading.

Gameplay

A second QoE is how well the FLAME platform is suited for this type of gameplay. We measure this QoE by understanding how engaging the game was for the trial participants and if the design of the FLAME platform prohibits certain gameplay elements. Overall, users picked up quickly on how to interact with the AR elements of the game, which were distributed around Millennium Square in Bristol. An introduction was given to the players shortly before the trial. During the more competitive games, we observed players running between the AR garden and gnome shops to get an advantage at the game. Users therefore covered a significant distance while playing and also stated that they were exhausted after two rounds of 15 minutes each.

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

On the negative side, the lack of different strategies in the game lead to the general opinion that the game is not interesting enough to play for a longer time. Therefore, new gameplay elements and features could have a positive impact on the overall experience.

We conclude that the FLAME platform was in no form inhibitive for designing mobile AR urban arena games with multiple players but instead very well-suited.

6.2.2.2 QoS

From a **QoS** perspective, we were interested in evaluating how the FLAME network load varies depending on the number of players participating in the game. Since the download of trees happens roughly at the same time from several devices, we expect the Opportunistic Multicasting of FLAME to have a positive effect on resources used and network load of our database service.

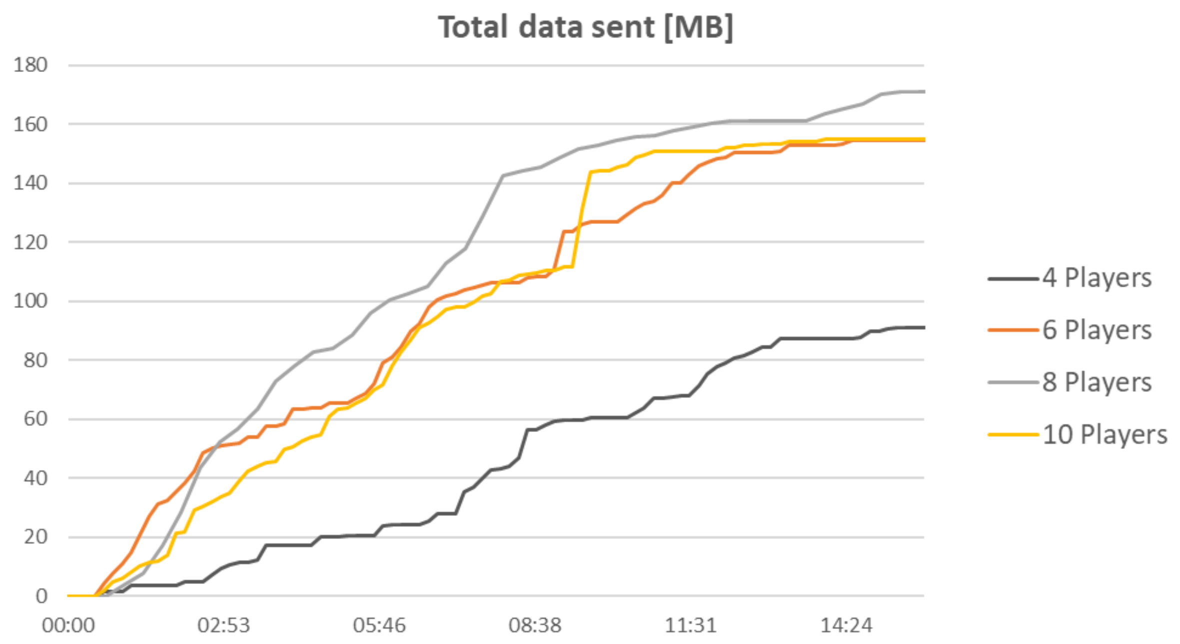

Figure 24: The total amount of data sent by the storage FMS over the time of a game for different number of players.

Network Load

We measured the total amount of data sent from our database service at a 20 second interval. Multiple plots for 4, 6, 8, and 10 players are depicted in Figure 24. Most noticeable is that in every game the curve reaches a plateau after some time. This is mainly due to the fact that the tests were performed with a limited number of 100 different tree models. Therefore, in the later parts of the game, newly planted trees were more likely to require a 3D model that had already been downloaded during an earlier stage of the game. In such cases, the game client did not have to request another download but simply load the tree model from memory.

Additionally, the data shows no significant difference in data downloaded between 6, 8 or 10 players, even though we would usually expect the data to increase linearly with the number of players. The same trend can be observed in the number of requests reaching the storage FMS, depicted in Figure 25. During an average game of Gnome Trader, we assume that most of the 100 tree models are planted and therefore also downloaded. With a normal server setup this would lead to approximately 100 x n

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

downloads, where n is the number of players. With FLAME, however, we never reached such high numbers, as can be seen in Figure 25. We also have to consider that a significant portion of the requests were due to players joining the game at a later point in time. For example, the 10-player game (yellow curve) had a total of 409 requests reaching the storage FMS. Approximately 55 of those were due to a player, who joined the game late and therefore had to download all currently planted trees at once. If there were no new players joining mid-game, we can conclude that the multicasting reduces the number of requests reaching the database by at least a factor of 2.

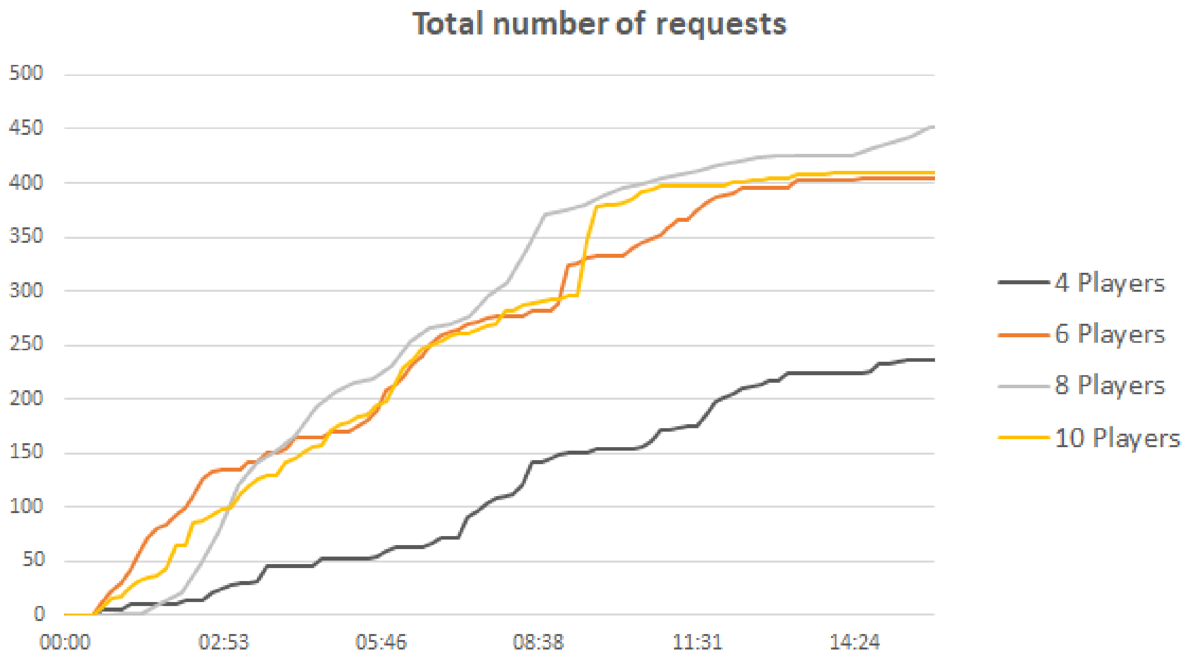

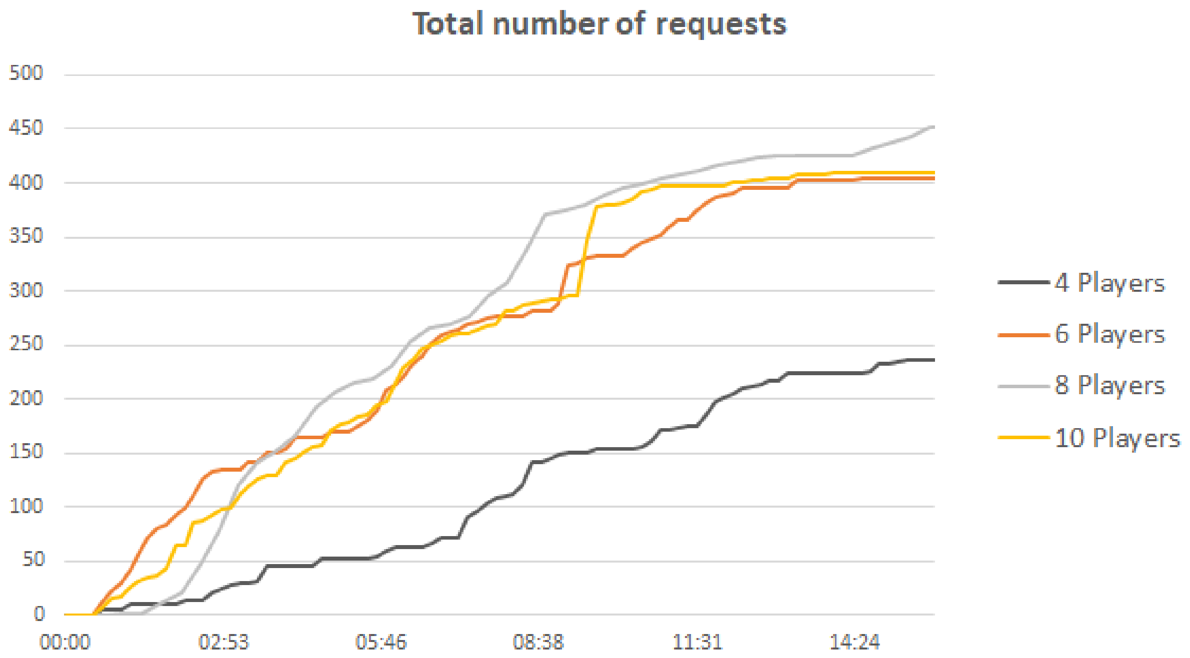

Figure 25: The total number of requests sent to the storage FMS over the time of a game for different number of players.

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

In conclusion, we measured that the FLAME platform significantly improves the network load QoS for our game by offering Opportunistic Multicasting.

Request Processing Time

Overall request processing time of the FMS storage service (Figure 26) is low except for moments when a new player joins the game. The player joining in the 10-player game lead to approximately 35 seconds of network activity. When many trees have to be downloaded, the client code is tuned to distribute the downloads over a longer period of time. In this trial, the 8-player game had the highest request processing time and also the highest number of players joining late or re-joining (3 occurrences).

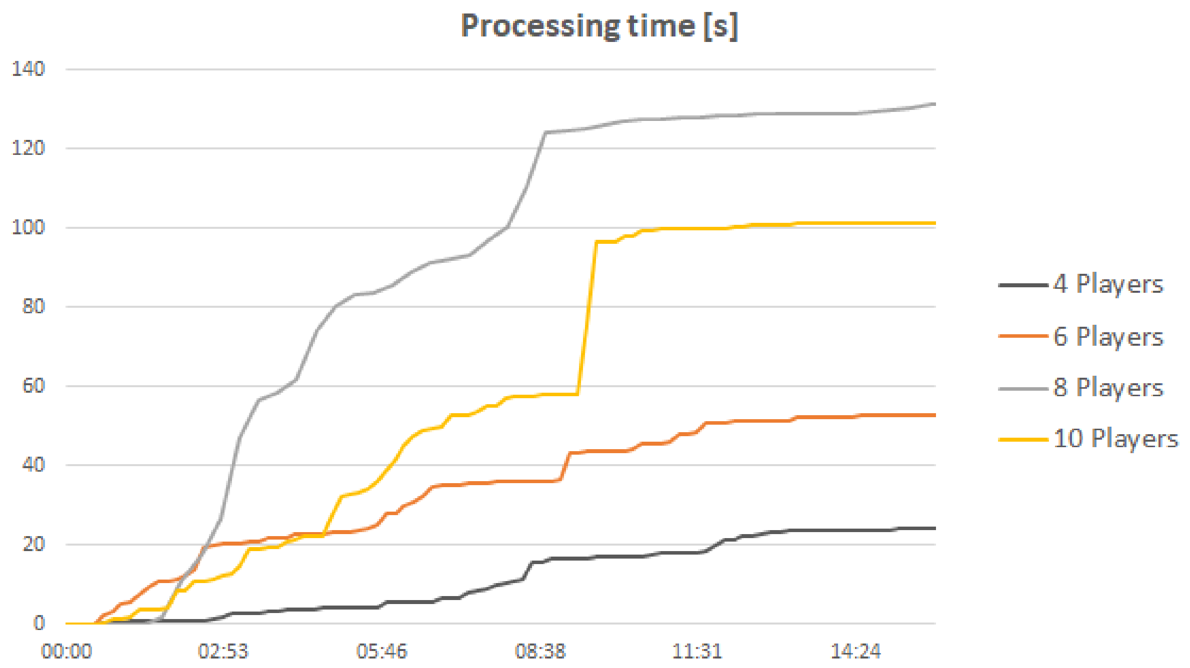

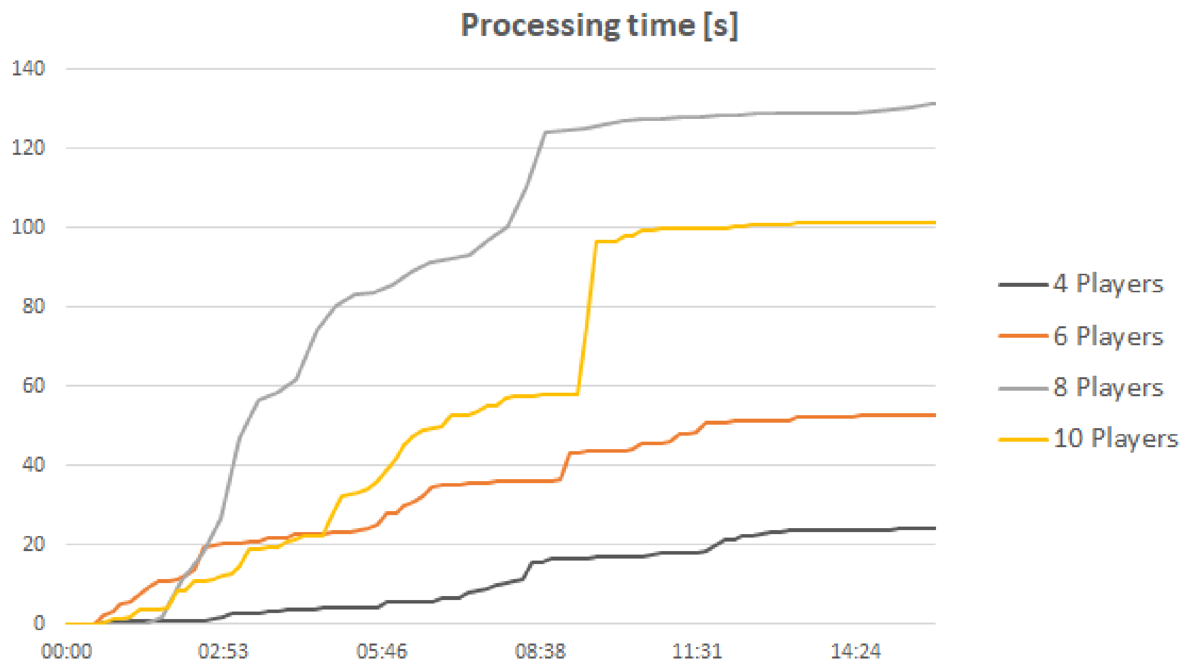

Figure 26: The network usage time of the storage FMS over the time of a game for different number of players.

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

The FLAME platform successfully handled all requests quickly and, even in games with 10 players, no long and gameplay-encumbering request processing times occurred.

6.3 INSIGHTS

Our evaluation indicates that the FLAME platform is capable of delivering 3D assets quickly and efficiently for video game applications. For larger number of players, the load on the server remained largely the same as for fewer players, which can be accredited to the Opportunistic Multicasting capabilities of the FLAME platform. To benefit from these capabilities, the underlying client code has to be tuned in a way that lets all clients request the same asset simultaneously. Players who were in the game since it started, never observed any delay in asset loading.

Players who join mid-game, which may happen in multiplayer games, put an over-proportional high stress on our test system. These were also the only players who observed a short delay while the 3D models were continuously loaded into the game. The main reason for this is that these players have to download all currently used 3D assets as soon as they join the game. This process takes time and the servers cannot profit from the Opportunistic Multicasting, as only one player is requesting these assets at that moment in time. However, as the content was served from the edge of the network, the problem was mitigated. This drawback is not a limitation of the FLAME platform but of the game design and implementation.

Finally, we conclude that the FLAME platform provides both high QoE and QoS as measured during our trials. Asset loading went smoothly and efficiently, while the gameplay remained uninhibited by the FLAME platform. Players were immersed in the virtual AR landscape and mostly unaware of 3D assets loading in the background. Opportunistic Multicasting helped at keeping the network load on our server at a minimum and request processing times were low, making the FLAME platform well-suited for urban multiplayer games. While this prototype of *Gnome Trader* targets the area of Millennium Square in Bristol, UK, and up to 10 players, future efforts could focus on exploring how similar games targeting larger areas, for instance, an entire city, perform on the FLAME platform. Our results indicate that such city-wide multiplayer games may be well-accommodated by the FLAME platform because of its scalability.

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

Expected outcome from experiment	Result
[Technical] Understanding if 3D asset distribution works over the FLAME platform	Achieved. Distribution of 3D models over the FLAME platform worked smoothly. 3D assets are quickly received by the client application and can be instantiated in the game.
[Technical] Understanding what effects Opportunistic Multicasting has on the server network load and how this design pattern can be exploited..	Achieved. In an average game, the number of requests reaching the server is approximately 2 times lower, than without multicasting. Still, in an average game there are many events that can circumvent the use of multicasting (e.g. player joining late, players having their devices locked, etc.), which makes it difficult to put an exact number on the performance boost.

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

7 CONCLUSION

This deliverable has presented the insights from broadcast, gaming and transmedia scenarios resulting from the validation experiments. Four pioneering FMI experiments were designed and implemented to explore the changing ways that consumers access and interact with localized media services. Each validation experimenter has described important insights both in the context of their target application as well as in the context of leveraging the FLAME platform to deliver new media services. An important theme in these insights is that the FLAME platform has become a usable and intuitive platform for delivering localized media services.

These validation scenarios were designed and developed during the early development of the FLAME platform. As such, they each demonstrate specific capabilities enabled by the FLAME platform. Each of the experiments have described technical insights of their applications but in addition the four validation experiments have helped drive the development of the FLAME platform and processes.

The validation experiments provided the first concrete requirements for the platform with initial brainstorming at FLAME General Assemblies resulting in ideas on what needed to be measured, what should be possible to change based on those measurements and how compute and content resources should be deployed and managed: informing the technical roadmaps for the CLMC and FLIPS components.

An understanding of the common media capabilities required helped inform the FMS roadmap, with several components initially tested with foundation experiments.

The first versions of the platform were validated by the four experiments, and not only technically: as their work progressed mistakes were made and rectified with changes to the methodology or the introduction of better tools and documentation.

In this way, the validation experiments have been invaluable in helping the project understand what was required and have provided generous feedback on the early technical and methodological aspects, helping bring FLAME to where it is today. The FLAME platform today is not only successfully utilized by validation partners, but also by 3rd Party Open Call participants. This further demonstrates the successes that can be achieved when media service provides, platform providers and infrastructure providers work together for the delivery of novel media services.

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

REFERENCES

[1] Dirk Trossen, et al. FLAME Deliverable D2.4. "Exploitation Roadmap and Action Report v1". 30/04/2019.

[2] Steven Poulakos, et al. FLAME Deliverable D5.2: "Large-Scale FMI Experiment Design in Broadcast, Gaming and Transmedia", 17/04/2018.

[3] Dirk Trossen, Michael Boniface, Gino Carrozzo. "Enabling 5G with FLAME". 31/05/2019. Available here: https://www.ict-flame.eu/publications/

[4] Mikolaj Leszczuk, Mateusz Hanusiak, Mylène C. Q. Farias, Emmanuel Wyckens, George Heston "Recent developments in visual quality monitoring by key performance indicators" Springer, Multimedia Tools and Applications, ISSN: 1380-7501 (Print) 1573-7721 (Online), pages 1-23, 6 Sept 2014

[5] Piotr Romaniak, Lucjan Janowski, Mikolaj Leszczuk, Zdzislaw Papir "PERCEPTUAL QUALITY ASSESSMENT FOR H.264/AVC COMPRESSION" IEEE CONSUMER COMMUNICATIONS AND NETWORKING CONFERENCE (CCNC), 2012

[6] Mu Mu, Piotr Romaniak, Andreas Mauthe, Mikolaj Leszczuk, Lucjan Janowski, Eduardo Cerqueira "Framework for the integrated video quality assessment" Multimedia Tools and Applications; ISSN 1380-7501, vol. 61 iss. 3, s. 787-817, 2012

[7] Eduardo Cerqueira, Mikolaj Leszczuk, Sherali Zeadally, Marilia Curado, Andreas Mauthe "Guest editorial special issue on "Future multimedia networking" MULTIMEDIA TOOLS AND APPLICATIONS, Sep 2011

[8] Mikolaj Leszczuk. "Assessing Task-Based Video Quality - A Journey from Subjective Psycho-Physical Experiments to Objective Quality Models" MULTIMEDIA COMMUNICATIONS, SERVICES, AND SECURITY, 2011

[9] Lucjan Janowski and Zdzislaw Papir. "Modeling Subjective Tests of Quality of Experience with a Generalized Linear Model" QOMEX: 2009 INTERNATIONAL WORKSHOP ON QUALITY OF MULTIMEDIA EXPERIENCE, Jul 2009

[10] Dirk Trossen, et al. FLAME Deliverable D3.10. "FLAME Platform Architecture and Infrastructure Specification v2", 11/01/2019.

[11] Michael Boniface, et al. FLAME Deliverable D5.6. "FLAME Ecosystem Operation Reports v1", 7/11/2019.

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