Operating a Next Generation Media Urban Testbed

A. Betzler*, C. Fernandez*, M. Catalan*, P. S. Khodashenas*, M. Lamarca†, S. Robitzsch‡ and M. Boniface§

*Fundació i2CAT; Barcelona, Spain (email: august.betzler@i2cat.net)
†Institut Municipal d'Informàtica de Barcelona (IMI); Barcelona, Spain (email: mariano@ieee.org)
‡InterDigital Europe; London, UK (email: sebastian.robitzsch@interdigital.com)
§ University of Southampton; Southampton, UK (email: m.j.boniface@soton.ac.uk)

***Abstract*—The adoption of 5G promises low latency and high throughput, greater adaptability of highly distributed compute and storage infrastructure, as well as configuration and management of a wide range of network services. This paradigm shift in network operations and usage results in a tighter integration of the infrastructure with the platform that allows instantiating the media services. The setup of the infrastructure in use for an urban testbed requires operators to carefully plan to acquire equipment with the expected technical characteristics, as well as to meaningfully configure, integrate with the platform on top and setup tools to monitor it. In this paper we share our experience of deploying an infrastructure on the street that supports the FLAME platform. Specifically, we introduce key points to consider when first defining the infrastructure site and equipment, when configuring, operating and monitoring it; and when experiments are to be tested.**

I. Introduction

The networked society is increasingly dependent on interactive multimedia systems. Today many systems are based on Over-The-Top (OTT) content distribution approaches where the delivery of media over the Internet is predominantly achieved through dedicated Content Delivery Networks (CDNs), thanks to the in-network placement of dedicated storage and network resources and without the involvement of a network operator. Such placement is facilitated by virtualisation and Software-Defined Networking (SDN), which thanks to its programmability, has fostered over the last decade the transformation of networking an computing models such as fog and mobile edge computing [1]; and also to Network Function Virtualisation (NFV), which places virtualised services at different points in the network. SDN has enabled application controllers to dynamically control topology and Quality of Service (QoS) [2] to improve observability of network traffic, as it has also allowed content isolation through slices and supported end-to-end communication channels over multiple mediums to distributed data centres [3]. Real-time requirements for media distribution have been translated dynamically to resource specifications in operator's clouds, including QoS characteristics like bandwidth, latency, packet loss and jitter.

In addition to the increasing softwarization of infrastructure, the content formats, consumption and production patterns are continuously changing as users demand improved Quality of Experience (QoE) and optimised media content delivery. This trend [4] expects four main key characteristics (PIML) to be addressed: personalisation, interaction, mobility and localisation. That requires adapting content to meet individual needs (*personalisation*) or to segments of users, often with a geographic component (*localisation*); as well as systems adapting to continuous and variable user input (*interactivity*) or changes on its location (*mobility*) and appropriately and timely reacting to them).

In this paper we describe the integration of the FLAME platform (which addresses PIML through custom routing, monitoring and slicing techniques) with the infrastructure on the street; resulting in an urban testbed for media experimentation. In the next sections we introduce different considerations that must be addressed for a successful interaction between the platform and the infrastructure and for infrastructure operation; as well as insights provided by our validation experiments.

II. The Platform

The FLAME project promotes a content-delivery solution that addresses emerging demand trends through cross-layer integration between Virtualised Service Networks (VSNs) and Information-Centric Networking (ICN) [5]. Virtualisation technologies expose virtual resources to functions in the application level, called Virtual Network Functions (VNF). In this context, shared HW resources are provided, thanks to technologies like SDN and NFV, in an isolated manner within the so-called slice of the infrastructure. Media services then run on top of the platform, isolated from each other.

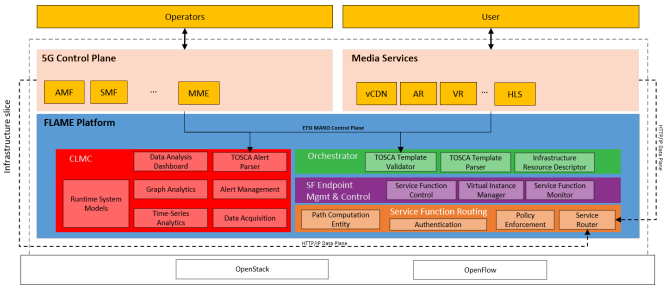

Fig. 1: Architecture of the platform

The FLAME platform (Fig. 1) provides the following components, each covering an aspect of the media service deployment: *Service Function Routing (SFR)*, *Service Function Endpoint Management and Control (SFEMC)*, *Orchestration* and *Cross-Layer Management and Control (CLMC)*. The SFEMC registers Fully-Qualified Domain Names (FQDN)

towards the SFR, which allows the service requests to be router to such Service Function Endpoint (SFE). For routing, an OpenFlow interface is used to insert suitable forwarding rules in the switching fabric of the underlying infrastructure; while the data plane Layer 2 Ethernet forwards traffic for both media service data and platform internal traffic. The Orchestration component interfaces with a custom resource management scheme that operates under the constraints of the underlying HW resources; which are specially limited in edge deployments. It exposes resources to media service providers in an optimised manner to ensure execution of the platform functions underneath. On the other hand, this component talks to SFEMC to realise orchestration-level management. Finally, the CLMC gathers information across layers and analyses the content flow in the network, as it may be needed for control-level decisions (i.e., reacting and adapting to certain conditions in the network).

The media services run on top of the platform, using its components for deployment, operation and monitoring. On the bottom lies the infrastructure, exposing an OpenStack-compliant interface so that the platform can operate with all virtual resources and offer these to the media services. The relation between platform and infrastructure(s) is crucial. A multi Point-of-Presence (PoP) can be used to expose resources from a distributed urban infrastructure in the same manner as a single PoP, yet allow to claim resources based on specific geospatial constraints, e.g. near the users.

III. The Infrastructure

We focus here on the main factors and physical constraints to consider before deploying the expected set of resources in an urban deployment. Therefore, installations should be planned based on desired concepts like coverage, availability and performance; which will dictate how and where the chosen equipment should be placed. Such a city setting consists of at least two sites: the on-street deployment that provides Radio Access Networks (RAN) capabilities and the Multi-Access Edge Computing (MEC) to provide light services close to the edge; as well as the main datacentre (DC). The sites are interconnected by an intermediate site and private networks.

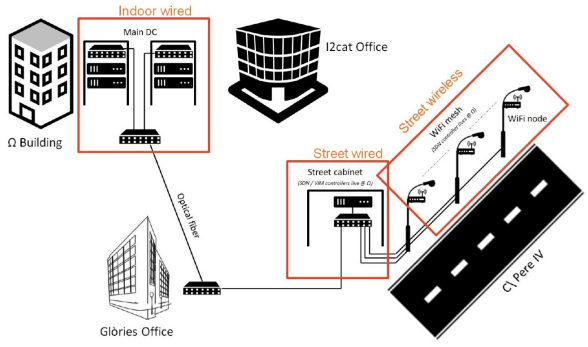

Fig. 2: Sites in the FLAME Barcelona infrastructure

Fig. 2 shows the high-level overview of the deployment in Barcelona, which will be used as reference throughout the section.

Beyond everyday consumer-oriented RAN technologies, many other types of technologies can be integrated in the Next Generation Internet urban deployment, such as sensors and actuators for smart city testing. We will though focus on the minimum required infrastructure.

A. Considerations on equipment

The infrastructure should provide the following types of equipment [7]:

1) ***Equipment with RAN capabilities****:* independently from the chosen technology, the wireless equipment is to be ruggedised for outdoors deployment (e.g., lampposts), there must be an easy way to access it and remote access or reset capabilities are recommended. Disregarding the installation spot of the radio equipment, it should be in a location that is secure, easy to access and maintain. There are also power, aesthetic and weight considerations to mounting equipment on street furniture which are usually defined by the manufacturers' specifications.

For our setup, we chose to mount each Wi-Fi node on a lamppost, acting as Access Points (AP), to provide connectivity for User Equipment (UE) in a pedestrian area. Within the chosen street (Pere IV, in the opposite side of the city from the main DC), a segment of around 400-500m serves the deployment of the Access Points (APs) that provide RAN capabilities. The lampposts host the mains power and fibre connections for the wireless nodes and were picked to be nearly equidistant to each other. They are connected via optical fibre with the FLAME edge (MEC) infrastructure. For the deployment of the wireless nodes, a third party designed a casing that fulfils the following requirements:

- Weather-resistant.
- Capable of switching from electrical to optical networks.
- Capable of converting from 220 V mains power (power line connectivity for standard "household" devices) to 48 V to power the wireless nodes.
- Providing a module that allows remote (hard) reset and an SNMP-based alarm system.
- Providing a battery that activates in case of loss of mains power, so the equipment can be turned off safely.
- Providing fans for ventilation, to keep the temperature in the casing below any critical threshold.

Inside the casing, a Gateworks Ventana (GW) 5410 Single Board Computer (SBC) [8], hosts several wireless network interfaces implementing the IEEE 802.11ac standard with backwards compatibility for the IEEE 802.11a/g/n standards. One interface is always used to instantiate wireless APs, whereas 1 or 2 additional interfaces enable optional wireless backhaul connectivity from each lamppost to its neighbours, providing experimenters with alternative network topologies to play with. For the RAN, omnidirectional dipole antennas are used; while for the backhaul directive panel antennas are

used (both supporting 2x2 MIMO). Using Ethernet, the GW SBCs connect over the optoelectronic media converter to the fibre leading to the edge cabinet.

The access of the UEs towards the platform must be evaluated when planning the RAN deployment for a city, since the type of RAN technology to use also depends on the experiment requirements. Most common technologies are Wi-Fi and LTE/4G, supported by the vast majority of smartphones or handheld devices (tablets/laptops). Other technologies, such as LORA, IEEE 802.15.4 or Bluetooth can be considered options for Internet of Things (IoT)-oriented scenarios, enabling connectivity for sensors, constrained devices and wearable devices.

*2) **Equipment with MEC capabilities**:* the MEC node or edge server is usually a shared resource, coming from the infrastructure provider, which runs light services close to the edge. It offers application developers and content providers cloud-computing capabilities on the edge, close to the end-user; who should benefit from improved QoE. For instance, the edge may support video analytic applications, location services, IoT, augmented reality applications, optimised local content distribution and data caching. Inside the MEC node, the Virtualised Infrastructure Manager (VIM) manages the virtual resources, in a similar way a private cloud provider would do. OpenStack is such an example. These virtual nodes enable operation of the platform and to run the experimenter's media services. Therefore, a degree of isolation (through the concept of slicing) must be achieved between resources that run in shared physical nodes but perform different tasks or experiments. In the VIM, this is typically achieved by creating different layers or overlays, each with a given range of VLANs; where such identifiers are used to delimit the scope of operation for each experiment and isolate the network traffic. In OpenStack, projects can be used to isolate resources and to connect to other virtual or physical resources, through the configuration of virtual or physical network devices.

In Barcelona, the MEC is a 12-core multi-threading CPU mini-tower server with 128 GB RAM and 2 TB of storage capacity. This machine acts as a compute node that is registered in the OpenStack controller, which is hosted in the main DC. Also deployed within the street cabinet there is a router (Cisco ASR920) that connects the equipment in the cabinet with the APs and with the main DC. It is configured, using VLANs, to properly forward specific kinds of traffic from and to the APs and the DC. It also provides L2 and L3 VPN services and offers high throughput and low power consumption

Similarly as in the DC, part of the resources in the MEC are allocated for the FLAME platform: in this case, one Service Router (SR) per lamppost. SRs are distributed at the ingress and egress of the network and are mandatory elements to realise the FLAME routing solution. The router provides enough Gigabit Small Form-factor Pluggable (SFP) ports to connect each fibre lamppost to the edge server, and other ports to connect the street with the main DC; where such connecting data and management communication lines are secured. The connection between the edge cabinet and the main DC has an intermediate hop in the IMI facilities.

*3) **DC equipment**:* main DC IT resources are used to provide heavy computational or storage services, e.g. high definition video content, video transcoding, quality of service and consumption analytics, as well as resource orchestration and management logic such as OpenStack, OpenDaylight, etc. The platform will likely require compute devices to host VNFs and should provide orchestration capabilities or allow talking to such software. Similarly, the platform is likely to be built on an SDN-enabled networking fabric. For instance, a VIM solution like OpenStack can be integrated with some NFV MANO as well as with SDN controllers.

The upper-left corner of Fig. 2 shows the Omega building that hosts the main DC infrastructure; where three servers are connected to each other following a star topology via a stack of two switches. That gives a degree of failure tolerance and High Availability (HA) in the computing cluster. As a VIM, OpenStack Ocata configured for self-service networks and Distributed Virtual Routing (DVR) with the Neutron OpenvSwitch agent is recommended. HA should be available as well for the project's virtual routers that reside on both the controller and the compute nodes.

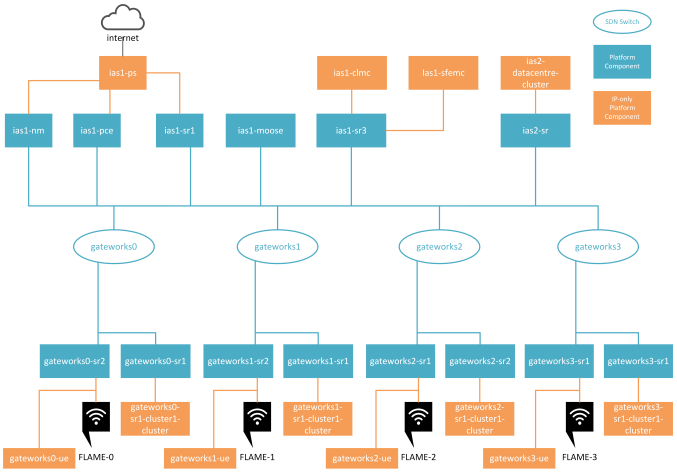

Fig. 3: Model of the infrastructure topology for the platform

Since the FLAME platform implements a stateless switching solution which requires the switches and controller(s) to be at least OpenFlow 1.3-compatible, a suitable SDN controller runs on the main DC. This controller must accept the rules via some API (e.g., REST) and insert them into the switches. The SDN controllers must support handling (read and insert) of arbitrary bitmask matching rules [6]. Floodlight and OpenDayLight controllers did support arbitrary bitmasks; while ONOS did not. We opted for Floodlight to make the SRs available to the platform. Fig. 3 shows the representation of the infrastructure architecture as understood by the platform; thus establishing a mapping between virtual SRs (defined within the platform) and physical GW APs (in the edge).

Regarding the SDN-enabled switches, there are considerations for either the silicon boxes and the software-based instances. The physical switches should support arbitrary bitmask matching via semantically overloaded IPv6 fields. As there is no capability verification alliance for OpenFlow and the OpenFlow 1.3 features are considered "experimental", it is highly recommended to double-check the specs of the device. On a software-based switching fabric, it is recommended to use OpenvSwitch. However, hardware switches implement the actual switch in their Ternary Content-Addressable Memory (TCAM) tables, which have an OpenFlow-compatible API and only one switch is known to support arbitrary bitmask matching, i.e. PICA8. The servers in the main DC are accessible over a fibre optic link from the i2CAT office and externally from the Internet.

4) ***Inter-site connectivity****:* different forms of connectivity should be considered depending on experiment needs: 1) fibre connectivity, which provides high data bandwidth. Installation is typically static and can be costly, as fibre has to be installed in ducts in the street and might need to cover long distances. Different topologies or combinations can be used, such as mesh, point to point, and star. Fibre is typically terminated in patch panels and a suitable media converter or SFP must be used; 2) microwave links, which give high throughput point-to-point connections and RAN connectivity for users. Links should be planned to take into account buildings and topography, typically delivering up to several hundreds of Mbps with IEEE 802.11. Installation cost is usually low, as physical installation is only needed at either end; 3) millimetre-wave links, which give point-to-point connections. Links should be planned to have line-of-sight and will typically provide up to 1 Gbps. Installation cost is usually low, as physical installation is only needed at either end; 4) fibre optic switches so as to interface between the backhaul technologies and access technologies. A switching fabric is required. Devices should support SDN to allow all devices to mesh together; support VLANs to allow traffic separation and be sized according to fibre or electrical split per location. Normally fibre is used in the ring and Ethernet towards edge devices. Required reliability and resilience can be provided by multiple switches or by more expensive switches.

The Barcelona FLAME infrastructure consists of the on-street deployment of the wireless nodes and MEC, and the DC (the Omega building) that are interconnected by a private network that goes through an intermediate site (the IMI premises). This site hosts networking equipment and acts as a concentration point of the fibre connections from the other sites (see bottom left in Fig. 2. The main DC and the MEC cabinet are connected through an optical network in two segments: 8 Gbps (DC to IMI) and of 8 x 10 Gbps (IMI to edge).

B. Considerations on management

When deploying on the street, all equipment is subject to local regulations. Fibre and radio infrastructure, equipment in the edge or even communication between sites. Specifications should be reviewed beforehand with the local authority to determine e.g., whether the equipment is properly adapted to the weather conditions (relative humidity %, supported operating and storage temperatures and possibly salt, dust or sand) and to the deployment site (power consumption, wireless power, distance between nodes and APs), etc.

Operation and Management (O&M) must be planned well in advance and contractors must be appointed according to expected timelines. Suitable budget must be allocated, considering that civil works can be costly. The preferred local government contractor is usually recommended as they will be familiar with the assets across the city. The level of service, and thus the celerity of response under failure, shall be adapted to the needs of the experiments.

Other possible regulations can affect which O&M works are allowed in the street and its timing and duration, or require notifications in time to local police in case of expected disruptions during works or experimentation. Finally, contracts and agreements may be needed between the infrastructure providers, the local authorities and the experiments to regulate terms of usage, data sharing agreements or others.

C. Considerations on operation

A key element when operating the infrastructure (and platform) during the trials was the continuous need for monitoring. A dedicated platform was setup to quickly show the status and notify about anomalies in the running services conforming the platform or in the on-street nodes and network. Any feature whose value remains too high (like latency) or too low (like the uptime) triggers a notification to operators about an apparent failure on any physical node.

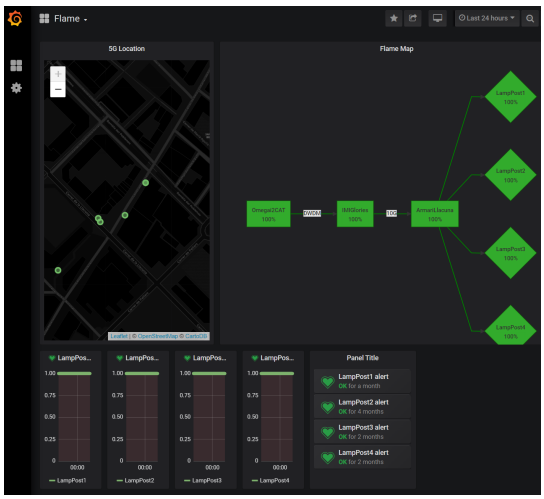

Fig. 4: Monitoring dashboard for the infrastructure

For our particular infrastructure monitoring setup, depicted in Fig. 4; a Telegraf server is used to gather SNMP data from the equipment (both Cisco ASRs and the Dell interface/VLAN statistics) along with InfluxDB to store it and a Grafana dashboard to render the statistics. Besides the physical monitoring,

the platform monitoring is recommended to clearly distinguish whether the problem comes from the physical nodes or from internal services running in the platform. Both monitoring systems should be checked periodically and especially right before and during a trial.

Another important aspect is security, both for the equipment and the services or data on them. Proper firewalls and access control should be in place at any equipment; specially those potentially reachable by users. Physical security is also a must in every node deployed in the street: these must be out of reach and protected by appropriate casing and measures to avoid unauthorised access on the equipment, such as "traps" and special ways to access the cabinet only known to operators.

IV. Validation after experimentation

Both infrastructure and platform were validated through user-centric trials on the street. Each trial leverages multiple interaction and service design patterns from the following: *opportunistic multicast*, *synchronised playout*, *nearest playout*, *proxy cache playout*, *content placement*, *application function offloading* and *scaling geographically* [9]. These patterns are enabled by the platform, which implements the logic to make these work; whereas the infrastructure lays out the physical foundation to support them.

The execution of these trials provided useful operational information that can be used to enhance both the platform and the infrastructure. Examples of such results are:

- Rainy weather conditions can significantly reduce the throughput on RAN transmission. The area taken by the experiment can also span some dead spots that have low coverage. Experiments should therefore be carried out ideally during non-rainy days and experimenters should check for areas with good natural (or extended) radio coverage; or otherwise be ready to properly tune (and reset afterwards) power settings in the APs.
- Streets with low density of nodes can be more easily impacted by external agents. An example of that is a high amount of traffic, which resulted in a decreased Signal-to-noise ratio (SNR). In such cases, it is recommended to discard the current tests and retry afterwards.
- Configuring the basic settings for the APs (like transmission power) should be as easily accessible as possible, and depending on the number on nodes and expected frequency of changes, also prone to automation; so that expected settings can be directly loaded prior to each experiment.
- Any non-experiment related data traffic can potentially harm an experimental platform. During the first validation tests and experiments it is recommended to switch off any possible source of external data traffic that could interfere with the logic of the internal routing of packets.
- Initial throughput tests indicate the maximum rate supported by the APs in the street. It is recommended to define the maximum threshold for parallel transmissions (e.g., in media streaming servers) allowed at once.

V. Conclusion

Deploying an urban infrastructure that is ready for efficient Next Generation Media transmissions is challenging for operators and can be a month-lasting task in the best scenario. In this paper, we compiled our experience of the definition and configuration processes in the form of key factors and considerations, so that the task of deploying similar infrastructures in the street is eased. We document as well the insights obtained after the validation process took place; whether through internal experiments and external trials. This provides further conclusions after validation that complement the guidelines described in the former sections.

Acknowledgment

The research leading to these results has been supported by the EU funded H2020 project FLAME (no. 731677) and Spanish national project ONOFRE-2 (no. TEC2017-84423-C3-1-P). Authors would like to thank all the involved colleagues on the design, deployment and operation on both infrastructure and platform.

References

[1] Hu, Y.C., Patel, M., Sabella, D., Sprecher, N. and Young, V., 2015. *"Mobile edge computing — A key technology towards 5G"*. ETSI white paper, 11(11), pp.1-16.

[2] Gorlatch, S., Humernbrum, T. and Glinka, F., 2014, February. *"Improving QoS in real-time internet applications: from best-effort to Software-Defined Networks"*. In Computing, Networking and Communications (ICNC), 2014 International Conference on (pp. 189-193). IEEE.

[3] Channegowda, M., Nejabati, R., Peng, S., Amaya, N., Zervas, G., Shu, Y., Rashidifard, M. and Simeonidou, D., 2013, September. *"Design and demonstration of multi-domain, multi-technology software defined networks for high-performance cloud computing infrastructure"*. In Optical Communication (ECOC 2013), 39th European Conference and Exhibition on (pp. 1-3). IET.

[4] Poulakos, S., Sumner, R., Crowle, S., Boniface, M., Phillips, S., Young, G., Matton, M., Carrozzo, G., Zuend, F., Trossen, D. and Garcia, J., 2017. FLAME D3. 1: FMI vision, use cases and scenarios. Available at https://www.ict-flame.eu/deliverables/

[5] M. Boniface, D. Trossen, and M. Calisti, *"User-Centric Media Demands through Software Defined Infrastructures"*, NEM Summit, Portugal. 23-25. Nov 2016.

[6] N. L. S. da Fonseca, R. Boutaba, *"Cloud Services, Networking, and Management"*, Wiley, 2015.

[7] FLAME D5.1: FLAME Replication Process v1. Available at https://www.ict-flame.eu/deliverables/

[8] Gateworks Ventana 5410 Datasheet. Available at http://www.gateworks.com/imx6-single-board-computers-gateworks-ventana-family/item/ventana-gw5410-network-processor

[9] Interaction and Service Design Patterns in FLAME, version 1.1. Available at https://ict-flame.eu/wp-content/uploads/sites/3/2019/07/FLAME-Service-Design-Patterns-v1.1.pdf