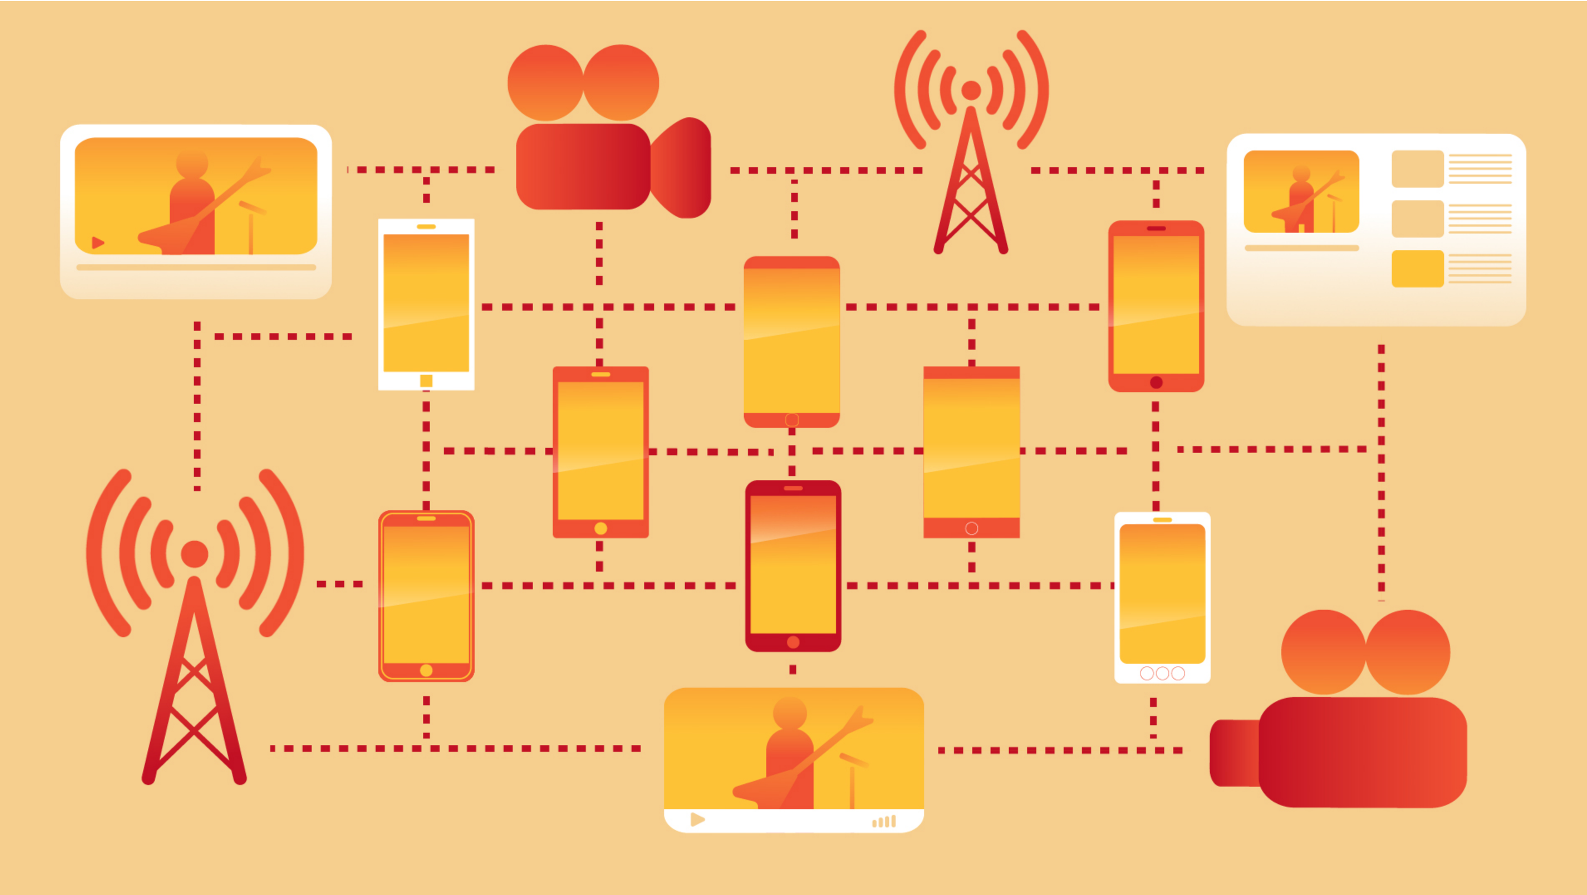

Enabling 5G with FLAME

Dirk Trossen (*InterDigital Europe*) | Michael Boniface (*University of Southampton*) | Gino Carrozzo (*Nextworks*)

31/05/2019

WWW.ICT-FLAME.EU

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

DISCLAIMER

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

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

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

EXECUTIVE SUMMARY

Mobile broadband communication has pervaded the communication market, responding to the most common and compelling nowadays users' needs for continued connectivity, fast browsing and audio/video streaming on handheld portable devices. Building on the 4G success through LTE and LTE-Advanced together with recent deployments in large-scale country-wide Operators' networks, **5G** is rapidly becoming a market reality. Various live technology trials worldwide are coming to completion, commercial networks in multiple countries and terminals are being announced to go live well before the "*ambitious*" target set by industry for year 2020. Consumers' expectations are the main spin for a fast move towards 5G [*ERI-REPORT*], due to needs for consistently improved network performance, reduced urban network congestion and wider more performant network coverage for mobile terminals of various kinds (e.g. handheld smartphones, drones, connected cars, robots, IoT sensor and actuators, etc.). The **core set of technologies** for 5G has fast matured during recent years to provide a new scalable network architecture which can meet exponentially increasing demands on mobile broadband access, both in terms of the number of connected users and advanced network-intensive applications.

The **Future Media Internet** has been evolving along these developments in mobile broadband with key technologies, such as edge computing, orchestration-based service deployment, and programmable infrastructures, being key to its realization. As a project operating at the heart of FMI, FLAME has therefore developed technology innovations that are key to 5G and FMI alike, as being notable in the standards contributions to key 5G forums, such as 3GPP and IETF, by FLAME partners. Furthermore, the urban media deployment propositions developed by FLAME partners, most notably Bristol and Barcelona, clearly align with expected 5G offerings such as Platform-as-a-Service provided by, e.g., smart cities via neutral host models.

This whitepaper is exploring the intersections between 5G and FMI from the specific angle of the FLAME contributions to it. For this, we will outline the main 5G drivers that also underlie the work in FLAME, the main contributions of FLAME to 5G developments and the roadmap to 5G impact through our FLAME activities.

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

1 5G DRIVERS

Technology advancements for 5G have been driven in last years by three key drivers, from which all the 5G service platforms and solutions eventually originated. Let us outline these key 5G drivers against which we align the benefits of the FLAME platform later.

1.1 PROGRAMMABLE INFRASTRUCTURE

5G realizes the concept of not just one-size-fits-all network architecture but many possible architectures for different verticals, aiming at different requirements as well as deployment footprints. In order to provide the necessary **flexibility**, the infrastructure in 5G is programmable in terms of connectivity and compute/storage resources. More specifically, technologies like SDN [*SDN*] and NFV [*NFV*] are being positioned to provide **programmable** connectivity between equally programmable network functions deployed deep in the network, utilizing **virtualization** of compute as well as connectivity resources. With this, not just commonly IP-based Internet services are targeted but also advanced low-latency, high bandwidth as well as localized 5G services. SDN-based solutions as well as virtualization platforms are not new and have been developed for several years preceding the development of 5G, particularly in the data centre community, specifically in cloud data centres. 5G adopts these solutions and extends their use beyond the data centers to implement flexibility also in other sections of the network, i.e. edge cloud, multi-access edge/far-edge, radio access, etc., to bring service dynamicity closer to the end-user.

1.2 CLOUD NATIVE DEPLOYMENT OF 5G

5G pushes programmability even further by adopting the cloud community insights into virtualization and SDN through what is called a **cloud-native deployment**. In other words, the infrastructure for 5G becomes a **highly distributed data centre**, made up of many micro data centres in regional offices and even at street level, all interconnected over high throughput programmable Layer2 so-called **SD-WAN** (software-defined wide area networking) connectivity. This infrastructure deployment model is complemented with work in 3GPP to shift 5G networks design from the pre-5G point-to-point design to a fully meshed **service-based architecture** (SBA) [*3GPP-SBA*], where network functions and user equipment interact based on the same service invocation interface. For the latter, HTTP as the most dominant protocol in the Internet has been chosen by 3GPP. Such SBA-based design will make the development of specialized control and user planes for future verticals akin to the development of **cloud-based applications** running over the same (HTTP) protocol, increasing flexibility for operators to provide specialized solutions while minimizing costs for doing so. The adoption of this design pattern is often captured under the tag line “**Telecom meets IT**”.

1.3 DEEP IN-NETWORK SERVICE DEPLOYMENT

5G utilizes programmable, cloud-native capabilities to enable **micro services** [*MSERVICE1*] being deployed deep in the network as well as the its edge. **Orchestration** frameworks, such as those developed in [*NFVMANO*], enable the deployment of such novel 5G services in a fraction of time compared to service deployment in available 4G and fixed networks. The **lifecycle management**

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

of those services will align with the dynamic nature experienced in cloud data centres, therefore bringing the **dynamicity** of the cloud to the entire 5G system.

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

2 KEY FLAME CONTRIBUTIONS TO 5G (AND BEYOND)

The FLAME platform architecture and technologies have been designed and implemented inspired by the same key drivers outlined in the previous section. As such, key FLAME technologies can bring benefits to 5G overall.

2.1 DRIVING ARCHITECTURE ALIGNMENT

FLAME has adopted a fully service-based architecture from the get-go, as specified in [*D3.10*] with HTTP being the main **service invocation** protocol alongside the entire Internet protocol suite. The FLAME platform furthermore integrates a **dynamic lifecycle management and control** (see also Section 2.3), combined with a novel, highly **dynamic service routing** (see Section 2.4) between the micro services (see Section 2.2) deployed over the FLAME platform and integrated with state-of-the-art **programmable transport infrastructures** that are 5G ready (see Section 2.5). With that, FLAME not only touches upon all the ambitions of a fully 5G-compliant SBA platform but it has also been at the forefront of demonstrating its capabilities in world's first public demonstrations in collaborations with 5G players such as Deutsche Telekom [*Release1*][*Release2*] at major industry forums such as the Next Generation Mobile Network (NGMN).

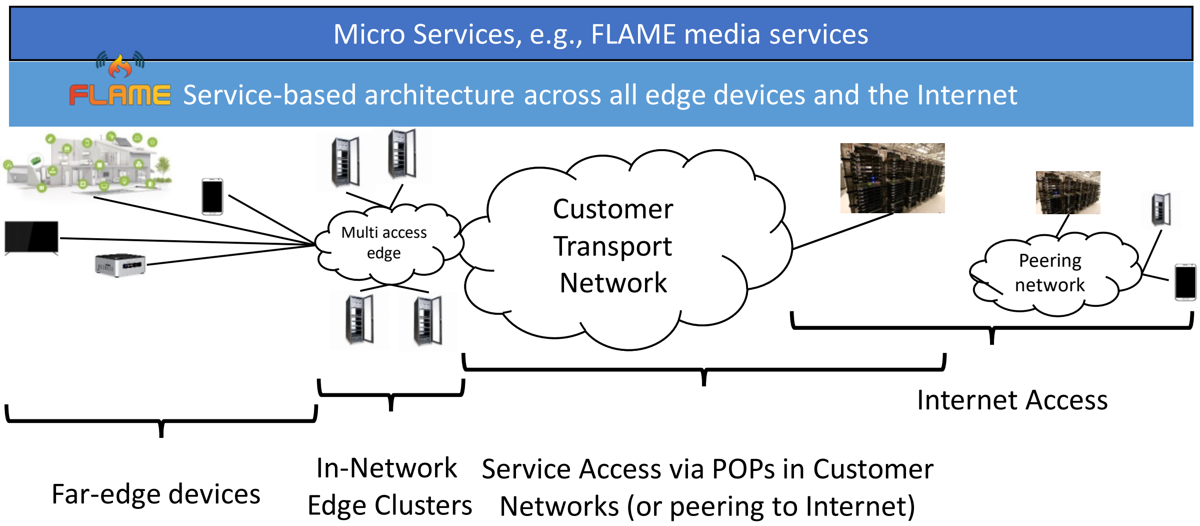

Figure 1: FLAME- Spanning from Distant to the Far Edge Cloud

But more importantly, FLAME has been driving the **architecture alignment** of the SBA model from the distant, often centralized, cloud to the far edge of the (in the future 5G) network, as illustrated in Figure 1. Early work in developing much of the underlying service routing technology (now finding entry into 3GPP and IETF standardization) covered utilizing the SBA model for the customer networks of existing and future (5G) operators, while early trials in 2017 [*Point17*] as well as current trials through FLAME address the suitability of the solutions for edge clusters found in deployments such as in Bristol and Barcelona pilots. More recent work, showcased at the Mobile World Congress in 2019 [*Release3*] extend the reach of the SBA concept to the mobile terminal

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

and beyond, positioning FLAME at the cutting edge of research in aligning micro service-based realizations of use cases with the same conceptual model, utilizing the same architectural foundations as well as utilizing key FLAME technologies in enabling them at scale.

Combined with its thrust into key SDOs, most notably 3GPP and IETF, FLAME has been at the forefront of the architecture alignment that embodies the 'Telecom meets IT' tag line mentioned in our 5G drivers overview.

2.2 SUPPORTING CLOUD-NATIVE SERVICE PATTERNS OUT OF THE BOX

The focus on the 5G vision of cloud-native deployment has also driven the contribution of FLAME towards service definition and orchestration. Great emphasis has been placed on supporting service patterns that are more aligned with such micro service over an SBA platform approach than modelling a mere virtualization-based service deployment. The latter approach has been driven by the NFV efforts [*NFV*], requiring the specification of virtual compute instances together with the virtualized connectivity interconnecting each instance. The mental model here is more aligned with the traditional telecom deployment through connecting physical machines albeit realized through virtualization.

As illustrated in Figure 2, **cloud-native service patterns**, on the contrary, postulate the establishment of **service function chains** that capture the **business logic** a customer wants to realize, while each element of the chain being realized as a (micro) **service function** which in turn can be realized through one or more **service function endpoints** that are being deployed in the network as part of the **service**. Said endpoints execute in (often) virtualized resource **instances** of the **infrastructure**. The invocation of service functions within the service function chain is realized through an intelligent **service function routing** mechanism, akin to the routing in a cloud data centre, natively supporting the distribution of more than one instances across the regional data centres. Said routing is provided by the aforementioned SBA platform currently being standardized in 3GPP for 5G systems, aligned with techniques utilized in cloud computing.

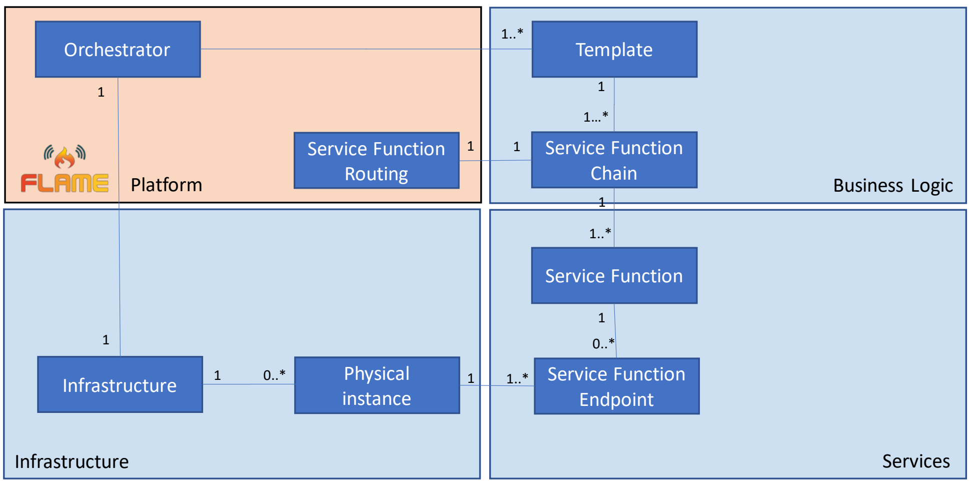

Figure 2: Supporting Cloud-Native Design Patterns in a 5G Environment

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

Through the service definition and orchestration solutions provided by the FLAME platform, cloud-native operation is being support out of the box, where the development insights have been informative for efforts in 3GPP and NGMN.

2.3 COMBINE ORCHESTRATION WITH CROSS-LAYER AWARENESS

Service-based systems are complex systems composed of many interconnected services and resources subject to unpredictable requests and affected by external events that cannot be controlled. In addition, systems are typically provided by multiple stakeholders cooperating through service offerings and provisioning of different types of resources across the full stack of distributed computing infrastructure, services and applications. This distribution and complexity make it difficult to have accurate and timely state information needed for orchestration decisions, especially when data is owned by different stakeholders.

FLAME's Cross-Layer Management and Control capabilities offer monitoring, measurement and analysis of service-based systems considering both temporal and dynamic topological characteristics of system elements contributing to performance. This creates **network-aware service function chains** that allow orchestration processes to understand how services respond to changes in workload and resourcing, and how such changes can be used to design, adapt and trial policies associated with Service Request Management, Fault Management and Configuration Management. The system provides mechanisms whereby service providers and platform providers can work together to improve the quality of shared state information as the basis for better resourcing and configuration decisions.

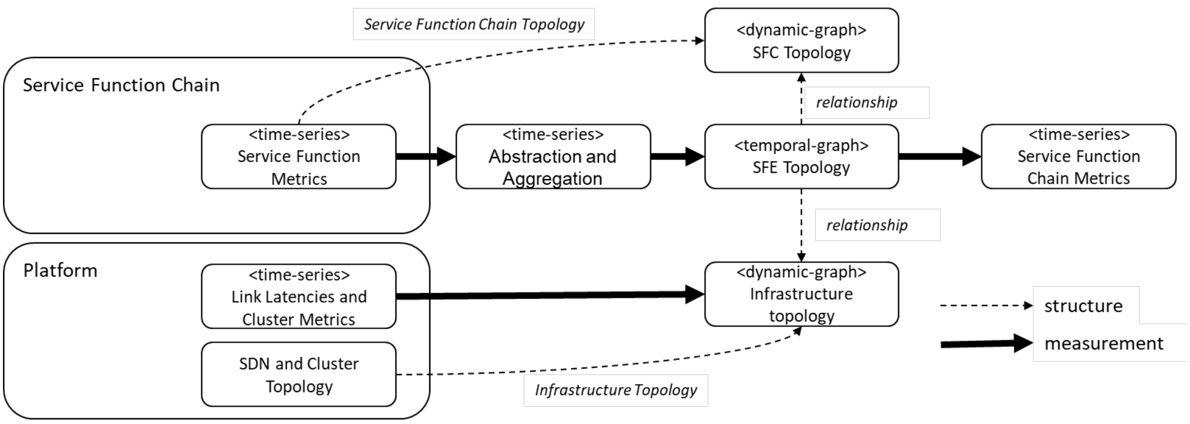

Figure 3: Cross-layer awareness

The measurement and analysis model is designed in relation to the lifecycle management stages of a service function chain (packaging, provisioning, management/control), with each stage in the process creating context for decisions, management and control in the next stage. Data Acquisition collects and hierarchically aggregates time-series measurements about system components from different layers including network and computing infrastructure, platform and services. The link between decisions and data is through queries and rules applied to contextual information stored with measurement values. Graph techniques are then used to relate time-

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

series data sets and allows complex analysis of system level properties, for example, end-to-end response time of a service function chain. The topology of a system graph is constructed dynamically from smaller topologies representing different aspects of the overall system (e.g. SDN topology, Service Function Chain, Service Endpoint Placement). Each topology has a different lifecycle and dynamic characteristics that need to be considered. Nodes represent configuration items with properties defining state derived from time services measurements, along with Edges between nodes defining relationships which themselves that also have properties (e.g. network latency between Clusters). The graphs are built dynamically and periodically to create temporal graphs used to sample higher level service function chain time-series measurements that are stored in the same way as low-level data points.

With time-series data available at the level of service function and service function chain, FLAME then allows Alerts to be defined that are used to trigger orchestration actions. Using Alerts service providers can specify how services should respond to changes in demand and resourcing in relation to desired Key Performance Indicators. Such actions could include managing the lifecycle or endpoints (e.g. place, boot, or connect), update routing/load balancing policy or trigger actions within services themselves to provide QoS-aware adaptation of content.

Such alert-based model allows for FLAME to provide an orchestration-based lifecycle management that provides reactive control level decision making with the underlying data for those decisions coming from all layers of the FLAME platform.

2.4 PROVIDE DATA-CENTRE LIKE SERVICE ROUTING

A key ingredient to the service-based architecture discussed in Section 2.1 and the realization of the cloud-native service pattern is the **routing of service requests** among virtualized service instances, distributed across possibly different geographic locations of contributing micro data centres. FLAME has been at the forefront of postulating the cloud-native deployment of 5G as one that assumes distributed computing resources **connected via Layer 2 connectivity** only. This is long reality for in-network core components, implemented in a few regional offices connected via SD-WAN connectivity.

But the solutions deployed in FLAME goes beyond this initial step and extend the Layer 2 connectivity not only to the user plane but down to the mobile terminal itself. FLAME partners have been leading in the development of so-called **5GLAN** [*3GPP_5GLAN*] solutions, which provide LAN-based connectivity to any mobile device within a 5G operator as well as in roaming use cases. With this, FLAME has extended the use of its novel SBA platform to any mobile terminal, while having already covered the core and edge network deployments, therefore closing the gap for realizing the architecture alignment outlined in Section 2.1. The resulting **end-to-end protocol stack**, including a simplified terminal stack (currently being implemented in FLAME) is shown in Figure 4, with the FLAME SFR (service function routing – see more in Section 3) component realizing the novel service routing.

The latest demonstration at the MWC 2019 showcased a world's first LAN-only based end-to-end SBA realization for a use case that saw a mobile application being dynamically offloaded between devices and in-network computing resources. This demo also utilized efforts in the IETF that have formulated the notion of **name-based chained transactions** [*Trossen18*], realizing the dynamic routing between virtualized instances of service functions through a novel information-centric

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

networking (ICN) based routing that allows for fast indirection between changing endpoint instances in the millisecond range. This brings the routing within the transport network close to the indirection latencies experienced in highly optimized cloud data centres, while truly supporting compute resources that span from the distant cloud to the terminal edge.

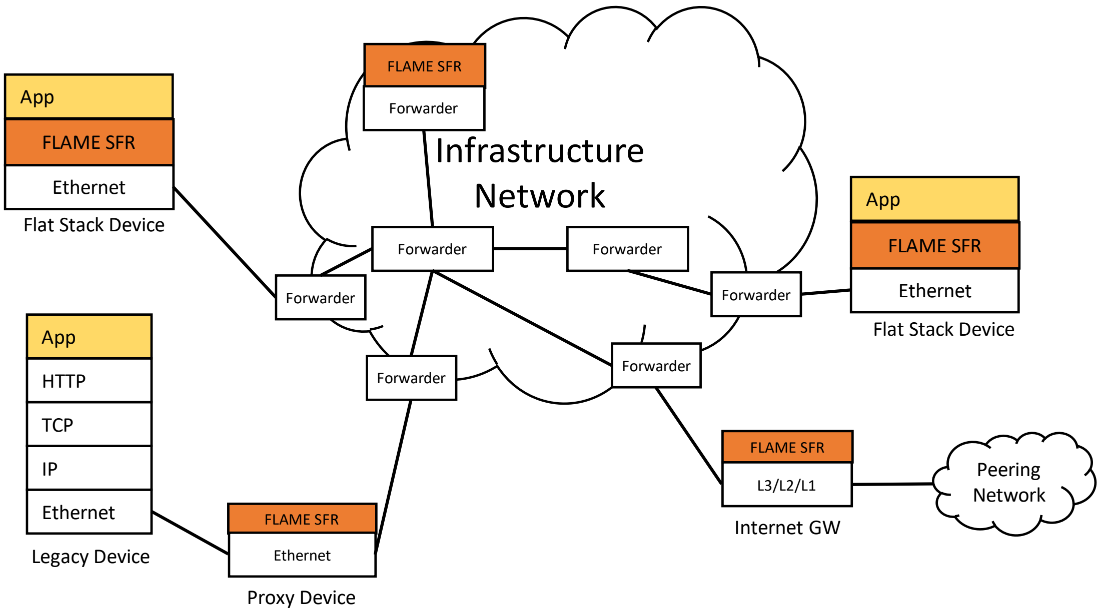

Figure 4: Provide Service Routing on top of Ethernet L2 Connectivity

Current contributions in this space have not only driven current 5G developments but will continue for upcoming 5G releases and beyond (contributions to the ITU-T Network2030[1] focus group can be found in [Trossen19]). This alignment places FLAME at the forefront of developing truly innovative end-to-end protocols for supporting micro services in highly distributed and dynamic environments.

2.5 INTEGRATION WITH 5G-READY SDN-BASED TRANSPORT NETWORK

The integration with a programmable transport infrastructure has been at the heart of the 'practical' work in the FLAME platform integration and deployment efforts. Together with the replicator sites in Bristol and Barcelona, **best practises** have been developed that allow for deploying a service platform such as FLAME in a network slice, utilizing state-of-the-art virtualization platforms such as **OpenStack** [*OpenStack*] and SDN control protocol like **OpenFlow** [*OpenFlow*]. These best practises are being documented and will be made available to the wider community through FLAME technical reports that accompany the platform documentation. This documentation will handle the specific aspects of **SDN** HW deployments**,** such as integration with

[1] https://www.itu.int/en/ITU-T/focusgroups/net2030/Pages/default.aspx

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

specific switches, but also the deployment over so-called whitebox, i.e., SW-based switches, covering general issues but also FLAME-specific aspects relating to the service routing and its novel path-based forwarding solution [*Reed16*]. Furthermore, FLAME will be engaging with the Open Infrastructure [*OpenStack*] community in particular for the **open sourcing** of useful configuration tools, such as the internally developed ARDENT toolchain, as well as deployment use cases.

Apart from integrating at the OpenStack and SDN level, FLAME has also deployed use cases with 5G radio support, such as through the UK-funded Smart Tourism [*5GST19*] project, in the Bristol deployment, which in turn has resulted in insights that will be documented for the wider community.

Through gathering the deployment and integration insights into concise best practise documentation and engaging with the operational community, FLAME is able to share its insights into real-life urban scale 5G deployments with the wider community and timely align this sharing with the trial and test bed efforts of the wider European and international 5G community.

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

3 COMING TOGETHER: 5G-READY FLAME PLATFORM

Figure 5 shows how our contributions to 5G come together in the developed, integrated and deployed FLAME platform, operating on top of an infrastructure such as the one provided in our deployments in Bristol and Barcelona. At the very bottom, we assume the existence of the **infrastructure (provider)**, exposing an ETSI MANO compliant interface [*NFV*] to the FLAME platform for resource management at the *wholesale level*, i.e. the FLAME platform is reserving platform resources in the compute, storage and networking domain. It is here where the best practises outlined in Section 2.5 are derived. We assume such infrastructure resource exposure provides the ability to reserve resources for the FLAME platform as part of a longer-lived relationship between FLAME platform provider and infrastructure provider, realized with an infrastructure slice.

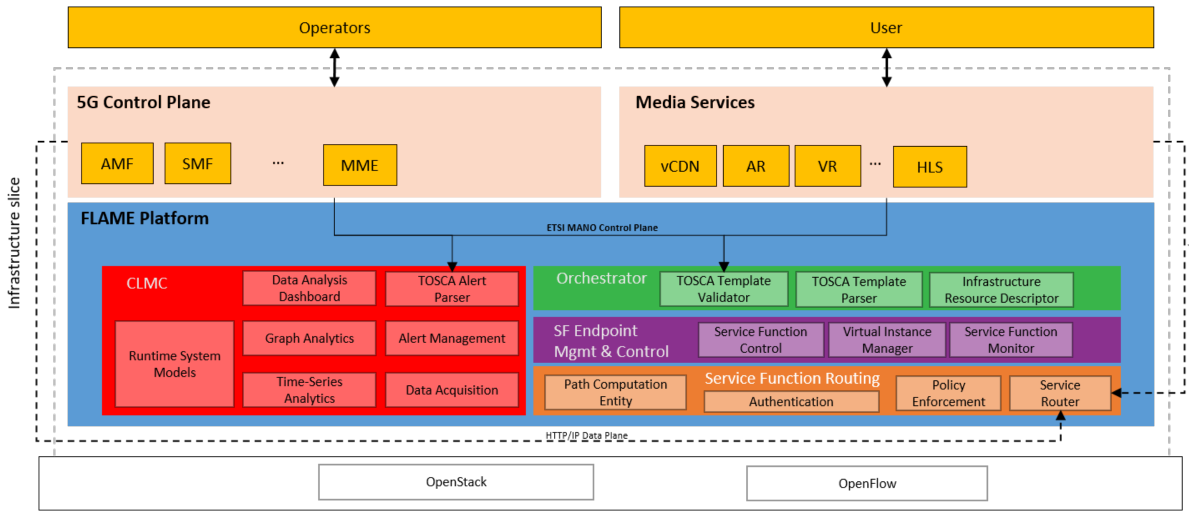

Figure 5: FLAME Platform Architecture

Resources of the infrastructure provided to the FLAME platform are in turn provided as *retail resources* to the (media) services (or to a 5G control plane) at the top of the platform. In addition, a monitoring interface allows for information exchange between application services and FLAME platform, which in turn will drive the decisions taken for the management input via the former interface. This combination of management and monitoring interfaces effectively provide a rich **demand-driven service deployment API** towards application services, allowing for defining and controlling the compute, storage and network resources, along with monitoring and alert configuration necessary to trigger service function policy actions and understand performance in relation to resource specifications, as briefly discussed in Section 2.3. For this, the management interfaces allow for initiating the orchestration of (retail) resources by a service provider. *We see these interfaces and therefore the dialogue between media services and the FLAME platform evolving over time.* In an initial realization, we see media services heavily relying on the ETSI MANO compliant interface, providing details on compute, storage and network resources being utilized in their specific deployment through **an explicit template**, based on the micro service patterns outlined in Section 2.2. However, empirical insights obtained from the initial deployments will

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

drive the evolution of the FLAME-specific (second) management interface towards one that provides a **selection of orchestration templates** for specific use cases. Such templates will therefore remove the need for explicitly defining each orchestration template from scratch, while utilizing the orchestration interface to utilize this evolving template knowledge towards the FLAME platform. This evolution is based on the evolving knowledge through our initial use cases on how to best accommodate demand and supply for specific media services. ***Ultimately, we see the relationship between application services and platform as being realized by high-level KPI-driven service-level agreements (SLAs), expressed through the FLAME-specific management interface and realized through an evolving CLMC on the platform side.***

As indicated in Figure 5, we consider the realization of the **services** outside the scope of the FLAME platform itself. We do assume, however, that a service is realized through a set of **service components**, each communicating through an HTTP/IP-compliant data plane interface. Service components are implemented utilizing platform resources, such as servers, connectivity components (e.g., switches) and others, while also utilizing resources outside of the scope of FLAME, such as end user devices and Internet-of-thing components. It is up to the policy of the FLAME platform provider if resources are provided exclusively to a service provider or from a shared resource pool. In the latter case, service provider deployments could run concurrently in the system, while the former case provides an exclusive access to the resources for a single service provider. The governance and specificity of the deployment will likely drive these policies. For instance, exclusive usage of geographically defined resources might make certain service deployments mutually exclusive since appropriate resources for another deployment are simply not available.

The **orchestrator** component of the FLAME platform interfaces with the infrastructure resource management, through the ETSI MANO compliant interface. The orchestrator manages the compute/storage/network resources towards the service provider, while it utilizes the policy control interface towards the **Service Function (SF) endpoint management and control (SFEMC) component** to realize the orchestration-level management policies as well as to set suitable shorter-term control policies for service function endpoints. For the realization of the configured service function endpoint policies, the SFEMC component utilizes the FQDN (fully qualified domain name) registration interface to control the registration and deregistration of the service endpoints towards the **service function routing (SFR)** component. With this, the 'visibility' of said FQDN towards service requests being routed can be controlled. The SFR component, in turn, will use the OpenFlow interface to suitably configure the switching fabric of the underlying infrastructure as well as implement the novel service routing outlined in Section 2.4.

As discussed in Section 2.3, particular consideration is given in our platform to the gathering of information across various layers, realized by the **Cross-Layer Management and Control** (CLMC) component in order to drive the aforementioned evolution of the service deployment API. While such data is useful and needed for control-level decisions, such as the activation of service endpoints, it also provides a rich pool of data for service providers to develop insights into resource specifications, adjusting crucial longer-term strategies such as those for content placement or media adaptation, and dimensioning SLAs expected to govern B2B and B2C relationships. The CLMC brings together time-series and graph analytics to understand demand, resourcing and performance properties of media service function chains deployed within the FLAME platform. For a given orchestration, the infrastructure and service function nodes remain largely static whilst the deployment and state of service endpoint instances varies throughout the

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

lifecycle of the service function chain according to demand and policies. The CLMC can continuously produce **temporal graphs** that integrate state, performance and resource usage of system elements that can deliver insights such as end-to-end delay for protocols at OSI L6 such as HTTP. The time-series measurements are acquired from existing monitoring frameworks provided by the individual sub-systems (such as those provided directly by switching platforms, obtained through the OpenFlow interface towards the underlying infrastructure) or monitoring frameworks offered by service components themselves.

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

4 ROADMAP TO IMPACT ON 5G

Our contributions to 5G can be summarized in the roadmap shown in Figure 6 along several SDO activities, most notably the 3GPP as the main 5G normative body together with the IETF.

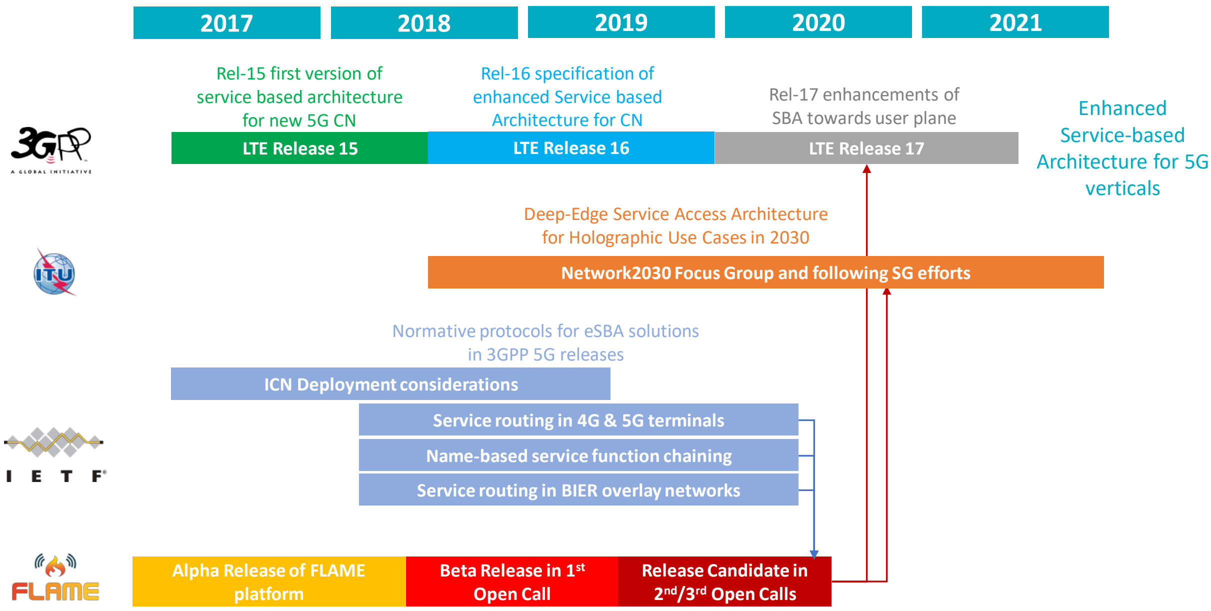

Figure 6: FLAME Roadmap and Impact to 5G-related Standards

FLAME started its impact on 5G in 2017, aligning early SBA contributions towards 3GPP Release 15 to establish micro service design patterns as the basis for 5G core networks beyond this initial 5G standard. During 2017 and parts of 2018, FLAME was working towards the **Alpha release** of its platform, while key technologies at the level of the SFR component (see Section 3) were brought to the IETF through the ICN (information-centric networking) research group and initial work on service routing for 4G and 5G in the same RG.

With the **Beta release** of the FLAME platform in time for the first open calls in 2018, contributions to 3GPP Release 16 and the IETF intensified with crucial contributions to the enhanced SBA and 5GLAN in 3GPP Release 16 and contributions to name-based service function chaining and Service routing in BIER overlay networks within the IETF. The availability of the FLAME platform supported those contributions through already mentioned key demonstrations to SDOs and forums such as NGMN.

Throughout 2019, while we continue normative contributions to 3GPP and IETF, we also expect insights from this normative work to find entry in the **Release candidate** for the FLAME platform, e.g., in terms of supporting service function chaining at the SFR component level. These insights will in turn be utilized to support future input into 3GPP Release 17 work, which is planned to commence in H2 of 2019, through the FLAME trial and demonstration activities.

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

Since late 2018, FLAME has also actively contributed to the ITU-T Network2030 longer term evolution efforts through contributions to service routing and edge architecture, positioning the impact of FLAME technologies not only for 5G but for future Beyond 5G, too.

As reported in [D2.4], FLAME has been a highly active contributor to 3GPP, IETF, ITU-T, ETSI MEC and NGMN with more than 70 contributions to these SDOs, many of which are placed in the normative solution space, positioning FLAME as a major contributor to the advancement of 5G!

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

REFERENCES

[*3GPP_SBA*] 3GPP, "New SID for Enhancements to the Service-Based 5G System Architecture", S2-182904, 2018

[*3GPP_5GLAN*] 3GPP, "5GS Enhanced support of Vertical and LAN Services". S2-186182, 2018

[*5GST19*] UK-funded 5G Smart Tourism project, at https://uk5g.org/discover/testbeds-and-trials/5g-smart-tourism/, 2019

[*D2.4*] "D2.4: Exploitation Roadmap and Action Report v1", FLAME deliverable D2.4, 2019

[*D3.10*] "D3.10: FLAME Platform Architecture and Infrastructure Specification V2", FLAME deliverable D3.10, 2018

[*ERIREPORT*] Ericsson Consumer & IndustryLab Insight Report, "5G consumer potential: Busting the myths around the value of 5G for consumers", May 2019, http://ericsson.com/consumerlab

[*MSERVICE1*] Dragoni, N., Giallorenzo, S., Lafuente, A. L., Mazzara, M., Montesi, F., Mustafin, R., & Safina, L. (2017). Microservices: yesterday,today, and tomorrow. In Present and Ulterior Software Engineering (pp. 195-216). Springer, Cham.

[*NFV*] ETSI, Network Functions Virtualisation, at http://www.etsi.org/technologies-clusters/technologies/nfv

[*NFVMANO*] ETSI GS NFV-MAN 001, "Network Functions Virtualization (NFV); Management and Orchestration", V1.1.1, at http://www.etsi.org/deliver/etsi_gs/NFV-MAN/001_099/001/01.01.01_60/gs_NFV-MAN001v010101p.pdf, 2014

[*OpenFlow*] OpenFlow, at https://www.opennetworking.org/sdn-resources/openflow

[*Openstack*] Openstack website, at https://www.openstack.org/

[*Point17*] POINT project, BIO Closed Trial, 2017, at https://www.point-h2020.eu/bio-closed-trial/

[*Reed16*] Martin J. Reed, Mays Al-Naday, Nikolaos Thomos, Dirk Trossen, George Petropoulos, Spiros Spirou, Stateless multicast switching in software defined networks, In proceedings of ICC 2016, Kuala Lumpur, Maylaysia, 2016

[*Release1*] InterDigital Inc., "InterDigital Demonstrates World's First Service-Based 5G Control Plane", at http://ir.interdigital.com/file/Index?KeyFile=394122639, 2018

[*Release2*] InterDigital Inc., "InterDigital Advances Cloud Native Deployment for Future 5G Wireless Service Delivery", at http://ir.interdigital.com/file/Index?KeyFile=395851857, 2018

[*Release3*] InterDigital Inc., "InterDigital Demos Comprehensive, Compelling Edge Visions at Mobile World Congress 2019", at http://ir.interdigital.com/file/Index?KeyFile=396854510, 2019

[*SDN*] Open Networking Foundation, Software-defined Networking (SDN) Definition, at https://www.opennetworking.org/sdn-resources/sdn-definition

[*Trossen18*] D. Trossen, D. Purkayastha, A. Rahman, "Name-Based Service Function Forwarder (nSFF) component within SFC framework", IETF draft, at https://tools.ietf.org/html/draft-trossen-sfc-name-based-sff-01, 2018

[*Trossen19*] D. Trossen, S. Robitzsch, S. Hergenhan, M. Reed, J. Riihijarvi, "Service Routing for Network2030", ITU-T Network2030 Focus Group contribution Net2030-SG3_054, 2019

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