

5G-Crosshaul Network Slicing

Enabling Multi-Tenancy in Mobile Transport Networks

Xi Li*, Ramon Casellas[†], Giada Landi[‡], Antonio de la Oliva[§], Xavier Costa-Pérez*, Andres Garcia-Saavedra*,
Thomas Deiß[¶], Luca Cominardi^{||}, Ricard Vilalta[†]

*NEC Laboratories Europe, [†]Centre Tecnològic de Telecomunicacions de Catalunya, [‡]Nextworks

[§] Universidad Carlos III de Madrid, [¶]Nokia, ^{||}Interdigital Europe

Abstract—5G requires a redesign of transport networks in order to feed the increasingly bandwidth hungry Radio Access Networks and to benefit of the performance/cost efficiency provided by the integration of both backhaul and fronthaul segments over the same transport substrate as well as the incorporation of Cloud RAN architectures. In addition, to increase its usage and cost-efficiency, this new transport network should allow simultaneous use by different tenants, e.g. MVNOs, OTTs, or vertical industries. This paper presents the 5G Transport Network architecture designed in the 5G-Crosshaul project to address this challenge. An SDN/NFV-based control plane has been designed that enables multi-tenancy through network slicing. The proposed solution allows for a flexible and efficient allocation of transport network resources (networking and computing) to multiple tenants by leveraging on widespread architectural frameworks for NFV (ETSI NFV) and SDN (e.g., Open Daylight and ONOS).

I. INTRODUCTION

5G mobile transport networks will support multiple Cloud RAN functional splits in a flexible and unified manner. This will allow for various degrees of Radio Access Network (RAN) centralization, varying from Distributed RAN (D-RAN) to fully Centralized RAN (C-RAN). Thus, 5G transport networks will flexibly distribute and move base station functions across data centres, introducing another degree of freedom for resource management. In this context, the division between *fronthaul*, which is the interface between the Remote Radio Heads (RRH) and their associated centralized-processing units (Base Band Units, BBU), and *backhaul* will blur, since varying portions of functionality of the base stations will be moved flexibly across the transport network, as required for cost-efficiency/performance reasons. In order to fulfill these requirements, we propose a new generation of transport networks for 5G integrating both fronthaul and backhaul segments into a common transport infrastructure, defined as **5G-Crosshaul** [1]. This 5G transport network aims to enable a flexible and software-defined re-configuration of all networking elements in a multi-tenant and service-oriented unified management environment, through unified data and control planes interconnecting distributed 5G radio access and core network functions, hosted on in-network cloud infrastructure.

One of the most important and desired features of 5G-Crosshaul is *multi-tenancy*, i.e. the ability to support multiple tenants while enabling flexible sharing of the 5G-Crosshaul physical infrastructure, so that each tenant can operate, independently, a subset of such resources. The aim of multi-

tenancy is to maximize the degree of utilization of infrastructure deployments and to minimize the costs of roll-out, operation and management – reducing both the capital (CAPEX) and operational (OPEX) expenditures – and to reduce energy consumption, which are essential goals of 5G [2]. In our context, a tenant can be associated to an administrative entity or user of a given service and implies a notion of ownership of one or more service instances and isolation between these instances.

Multi-tenancy is enabled by technologies such as network virtualization and network slicing, both covering the processes by which an infrastructure is physically or logically partitioned, segmented and assigned to different users. More formally, in line with related work, (e.g. [3]) we define a *network slice* as a *self-contained, coherent set of functions along with the infrastructure required to support such functions, offering one or more services for end-users*.

Although multi-tenancy is a concept that has been developed in many contexts, its applicability and benefits within transport networks has been addressed more recently. In the scope of 5GPPP, projects like 5G-NORMA ¹ or SESAME ² are addressing RAN multi-tenancy [4] while CHARISMA ³ covers 5G access networks. The work in this paper complements related work by focusing on the transport network aspects directly related to the combined fronthaul and backhaul, targeting per-tenant services which combine computing, storage, switching and transmission resource management. This paper presents a novel architecture unifying the aspects of resource virtualization, virtual infrastructure and network service management, combining the ETSI NFV Management and Network Orchestration (MANO) framework with integrated SDN-based control. *Note that the general concepts proposed in this paper can be also applied to other segments of a mobile system (namely the core and the RAN) to comprise an end-to-end (E2E) system. E2E network slicing relies on E2E orchestration (in some cases federation⁴) between different network domains.*

Our final target is to enable *Slicing as a Service* addressing the dynamic allocation of slices over a shared 5G-Crosshaul. The allocation of a slice involves the selection of the functions, their constrained placement, and the composition and config-

¹<https://5gnorma.5g-ppp.eu/>

²<http://www.sesame-h2020-5g-ppp.eu/>

³<http://www.charisma5g.eu/>

⁴<http://www.5gex.eu/>

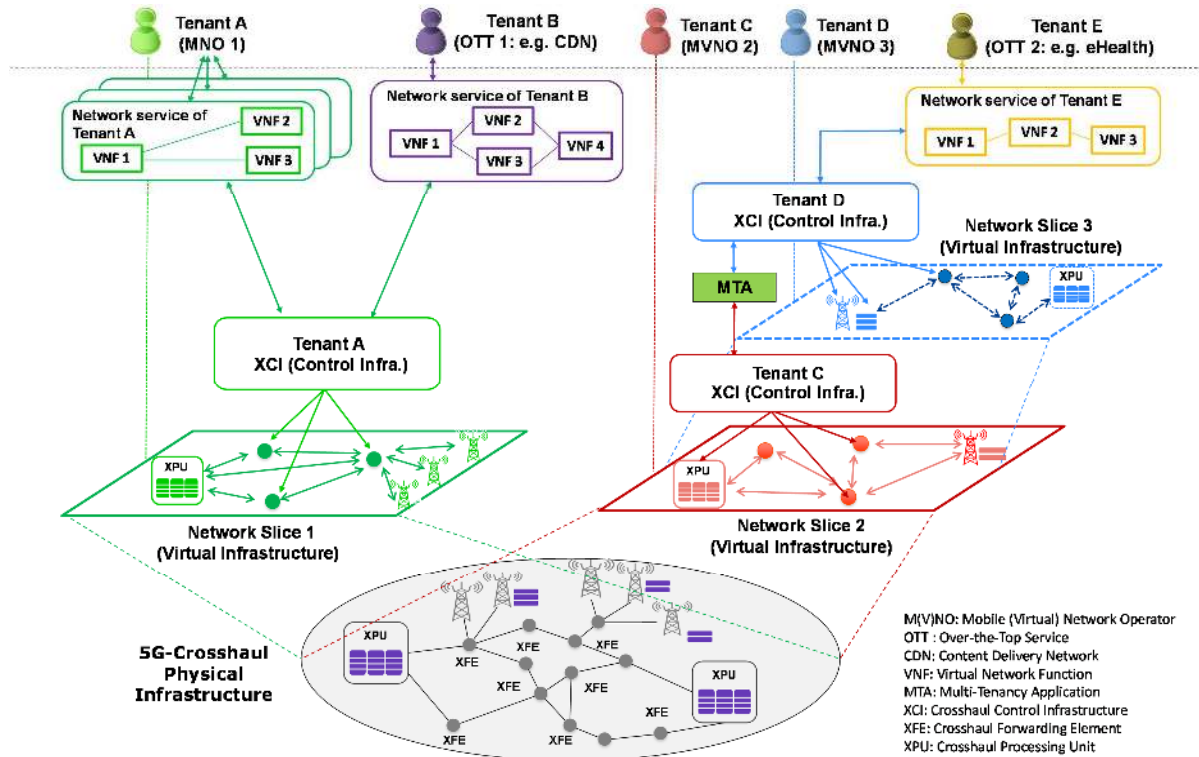


Fig. 1. Network Slicing in 5G-Crosshaul for multi-tenancy support

uration of the underlying infrastructures (either physical or virtual) fulfilling the services' requirements, in terms of e.g. latency, bandwidth or processing capacity. We consider two main network slicing services that enable different degrees of explicit control and are characterized by different levels of automation of network slices management:

- the provisioning of *Virtual Infrastructures (VI)* under the control and operation of different tenants – in line with an Infrastructure-as-a-Service (IaaS) model.
- the provisioning of tenant's owned *Network Services (NS)* as defined by the ETSI NFV architecture [5].

In the former, detailed in Section III, a *Virtual Infrastructure (VI)* is defined as a logical construct composed of virtual links and nodes, which, as a whole, "behaves as" and "can be operated-as" a physical infrastructure, enabling different degrees of internal control (i.e. can be operated by the tenant via different SDN control models). The service involves dynamic allocation of a VI, its operation and deallocation. The actual realization of a VI combines many aspects like partitioning and book-keeping of resources or the instantiation of connections supporting virtual links. The provisioning of a VI commonly requires direct hardware element support or its emulation via software for multiplexing over the shared infrastructure.

In the latter, described in Section IV, a NS is instantiated directly over a shared infrastructure, and as a set of interrelated

Virtual Network Functions (VNFs). A NS corresponds to a set of endpoints connected through one or more VNF Forwarding Graphs (VNF-FGs). Note that, whether the allocation of a NS is implemented in terms of the allocation of an underlying VI and the subsequent instantiation of the VNFs over the containing Virtual Machines (VM) is an implementation choice.

Multi-tenancy is an orthogonal characteristic of both services, guaranteeing separation, isolation and independence between different slices coupled with the efficient sharing of the underlying resources. Consequently, 5G-Crosshaul defines the term *Tenant* as a logical entity owning and operating either one or more VIs or one or more NS, ultimately controlling their life-cycle. The concept is illustrated in Fig. 1, where the owner of the physical infrastructure allocates virtual infrastructures over its substrate network, providing multiple network slices to offer to different tenants. Each tenant, e.g. a mobile (virtual) network operator (MNO or MVNO), owns and operates a network slice. In this example, tenant A, C and D owns the network slice 1, 2, and 3, respectively. Moreover, tenant A itself can also allow sharing of its infrastructure by other MVNOs. The MVNO tenants can further deploy their own NS or allow multiple third party tenants (e.g. Over-the-top service providers, OTT) to instantiate their NS on top of the virtual infrastructure, e.g. tenant B deploying its NS over the VI of

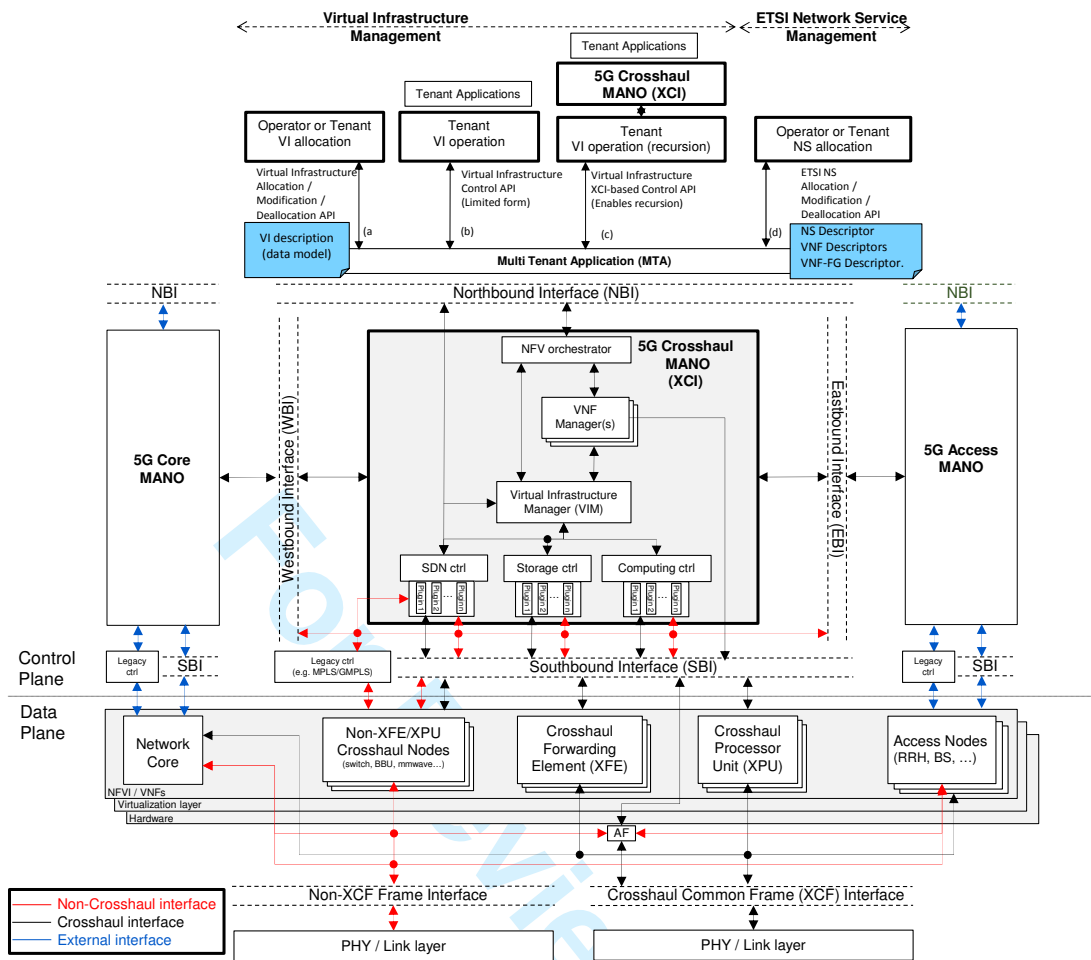


Fig. 2. 5G-Crosshaul architecture for multi-tenancy

tenant A. It is possible to instantiate a VI on top of another one following a recursive approach, e.g. the VI of tenant D is instantiated over the one of tenant C.

From the point of view of business models, network slicing allows MNOs to open their physical transport network infrastructure to the concurrent deployment of multiple logical self-contained networks. The availability of this vertical market multiplies the monetization opportunities of the network infrastructure as (i) new players may come into play (e.g. automotive industry, e-health, ...), and (ii) a higher infrastructure capacity utilization can be achieved by exploiting multiplexing gains. For the particular 5G-Crosshaul services, VI deployments are oriented to the B2B market, targeting customers like MVNOs or cloud providers specialized in customizable IaaS services, since they need a deep control on the network segment between distributed data centres. VIs can also be deployed by network operators to create virtualized and highly controlled environments to test and validate services before their roll out. Conversely, NSs target customers operating in the B2C segment, like application or service providers that offer services to end users (e.g. content providers specialized in streaming services).

II. 5G-CROSSHAUL ARCHITECTURE

The extended 5G-Crosshaul architecture [6], supporting several use cases of Multi-Tenancy, is depicted in Fig. 2. It follows the SDN principles (i) data and control plane are fully decoupled, (ii) control is logically centralized, and (iii) applications have an abstracted view of resources and states. Our design approach leverages state-of-the-art SDN and NFV architectures to maximize the compatibility and integration of the system design with the existing standard frameworks and reference specifications, and to allow the reuse of open source projects to facilitate its deployability while minimizing the implementation costs. The extensions we proposed on top of the baseline architecture are **the Multi-Tenancy Application (MTA) and a set of APIs to support the various multi-tenancy services**, as shown in Fig. 2, for the control of a VI or NS lifetime, instantiation, modification and deletion (API classes (a) and (d) in the Figure), and for the control of the VI in its limited or full-featured form (API classes (b) and (c), respectively).

The **data plane** is comprised of Crosshaul Forwarding Elements (XFEs) and Crosshaul Processing Units (XPU)s. XFEs

are switching units, based on packet or circuit technologies, that interconnect a broad set of link and PHY technologies using a common framing (Crosshaul Common Frame, XCF) to transport both backhaul and fronthaul traffic. XPU's take care of most of the computational burden including BBUs or MAC processors, VNFs and other virtualized services. To this aim, the data plane makes use of a NFV Infrastructure (NFVI) relying on generalized hardware components.

The **control plane** is divided into two layers: an application layer at the top and the 5G-Crosshaul Control Infrastructure (XCI) below. The XCI is our 5G transport MANO platform, compliant with the NFV MANO reference architecture, and provides an abstracted view of available resources, states and control and management functions to an ecosystem of applications, via a *Northbound Interface* (NBI). The XCI is connected to the data plane elements via a *Southbound Interfaces* (SBIs) to execute control and management functions on the actual hardware components. The NFV Orchestrator (NFVO) manages a NS life-cycle. It coordinates the VNFs life-cycle and the resources available at the NFVI in the data plane (supported by the Virtual Infrastructure Manager, VIM) to ensure an optimized allocation of the necessary resources and connectivity to provide the requested virtual network functionality. The VNF Managers (VNFM's) are responsible for the life-cycle management of VNF instances. Finally, the VIM is responsible for controlling and managing the NFVI computing (via a computing controllers), storage (via storage controllers) and network resources (via SDN controllers).

Although the scope of the XCI is limited to the transport network, it is essential to also consider the end-to-end coordination with other network segments (notably the 5G access and core segments). As shown in Fig. 2, our design includes a *West-/East-Bound Interfaces* (WBI/EBI) to communicate with the 5G Core MANO and the 5G Access MANO. They can be used for functions like reachability dissemination or (abstracted) topology and provisioning information to help achieve a system-wide optimization, enabling either a purely hierarchical architecture or a distributed/peer model for the orchestration of all involved segments. That said, the 5G access and core are out of scope of 5G-Crosshaul. Work in complementary projects like 5G-Exchange can be leveraged for multi-domain orchestration and federation [7].

The **MTA** is the application that implements the support for multi-tenancy, by coordinating and managing tenants' access to the shared infrastructure, driving resource allocation for instances assigned to different tenants, and delivering multi-tenancy related services by means of dedicated APIs⁵. A high level requirement is resource isolation, understood as the function of partitioning, separating and book-keeping of resources such that a tenant has no visibility of or access to the resources associated to another tenant. To perform this function, the MTA uniformly wraps and complements the infrastructure elements (e.g., SDN controllers, cloud management systems,

network elements, etc.) capabilities to provide multi-user and resource isolation support, offering uniform and abstracted views to tenants. Regarding mechanisms for isolation, our approach is to rely on existing ones, with the MTA acting as middle-ware and hypervisor. Full resource isolation requires system/infrastructure support and it is not straightforward or cannot even be achieved, e.g., without hardware redundancy. 5G-Crosshaul provides soft-resource isolation including, notably, driving the SDN controllers capabilities to create per-tenant networks, allocating software switches within XPU's dedicated to per-tenant traffic, defining security groups and per-tenant addressing, switching and routing within XPU's and logically separating traffic within XPE's. Similarly, from the ETSI NFV/MANO perspective, the MTA manages state regarding to allocation of Network Services mapping tenants to actual instances and relying on implementations support.

III. 5G-CROSSHAUL VIRTUAL INFRASTRUCTURE SERVICE

The allocation of a VI can be triggered by a tenant (such as a VNO), either directly consuming the MTA API - Fig. 2 API a) - or via the intervention of the infrastructure operator in a less dynamic environment, after an off-line Service Level Agreement (SLA). The VI concept is quite generic and can be extended to incorporate infrastructure elements beyond the ones considered herein. As part of the deployment of a VI, network, computing and storage resources need to be partitioned and aggregated, eventually recursively if a hierarchy is enabled. This partitioning can be committed in full at the time of instantiation (hard allocation) or reflected in terms of predefined quotas that are enforced at the time of use (soft allocation).

It is noteworthy that VI allocation follows an IaaS model, so the actual use of the VI (including the functions and related business logic) is defined by the tenant. The infrastructure owner is agnostic to the VI end use. Once a given VI has been allocated, the 5G-Crosshaul MTA empowers the tenants with different degrees of control to be exerted over it, with different operational models of control and management. In simple terms, this ranges between either: i) the control and management is restricted to the operational management and integration with tenant OSS/BSS and the operation of VI is mostly autonomous, with limited involvement of the tenant, such as monitoring and SLA validation, or ii) each tenant is free to deploy their choice of the infrastructure operating system and control plane, allowing the optimization of the resource usage within each VI. The former model involves the MTA offering an API that enables the tenant to have a limited form of control over the (abstracted) elements that constitute the VI - Fig. 2 API b) -, including a set of operations and policies that can be applied (e.g. retrieve an aggregated view of the virtual infrastructure topology and resource state and apply rules that affect element configuration and behavior). Low level operations such as the actual configuration and monitoring of individual flows at the nodes may not be allowed. The latter implies per-tenant controller - Fig. 2 API c) - or per-tenant MANO (XCI) including, most importantly,

⁵In the considered model, a single tenant entity owns one or more instances of each service in a 1:N relationship.

1 the ability to offer network services over its allocated virtual
 2 infrastructure. This approach ultimately enables recursion (as
 3 detailed in Section VI).
 4

5 IV. 5G-CROSSHAUL ETSI/NFV NETWORK SERVICES

6 The allocation of a NS extends and complements the
 7 concept of VI deployment – Fig. 2 API d –, to deliver isolated
 8 chains of virtual services composed of specific VNFs, in an
 9 automated manner. The tenant request usually specifies the
 10 type of VNFs (i.e. the desired virtual application components)
 11 in the NS Descriptor, their capabilities and dimensions through
 12 one or more VNF Descriptors and how they must be inter-
 13 connected through a VNF-FG Descriptor. Templates for the
 14 unified description of these information elements are currently
 15 under standardization process in the ETSI NFV ISG and in
 16 OASIS TOSCA standards [8].
 17

18 In the VI case, the tenant is responsible for the low-
 19 level deployment and configuration of its own applications
 20 over the allocated VI, while maintaining a certain level of
 21 control on the operation of the virtual resources. In the NS
 22 case, the tenant is interested in operating the applications that
 23 run in these virtual resources and expects that the needed
 24 level of resource capacity is seamlessly available in real-time
 25 without any further configuration effort. The deployment and
 26 continuous management of the whole service is completely
 27 automatized and totally delegated to the MTA and the NFVO
 28 within the XCI. The tenant has access to application-level
 29 interfaces only and the NS provisioning API follows an
 30 "intent-based" modeling approach where the tenant asks just
 31 for the composition of some network functions, without caring
 32 about how they should be deployed and delivered.
 33

34 In this scenario, the MTA is responsible for maintaining and
 35 coordinating the logical mapping between tenants, assigned
 36 services (in terms of NS and VNFs instances) and underlying
 37 virtual resources, in compliance with the SLAs established.
 38 Multi-tenancy is handled at different levels: at the lower level,
 39 a tenant has assigned physical and/or virtual resources in the
 40 domain of a Virtual Infrastructure Manager (VIM); at the
 41 upper levels, tenants have assigned VNFs and NSs. These
 42 different kinds of tenant can overlap and be merged in a single
 43 entity or be mapped over separate entities. For example a
 44 VNO can further virtualize the rented VI to serve different
 45 kinds of business customers, like CDN providers, delivering
 46 dedicated VNFs and NSs. The management of these tenants'
 47 relationships, together with the correlated authorization and
 48 SLA validation and assurance procedures are under the re-
 49 sponsibility of the MTA. Moreover, in these scenarios, NSs
 50 are not built directly on top of physical resources, but over
 51 Virtual Infrastructures through the allocation of VNFs and
 52 VNF-FGs in VMs and virtual network nodes, following a
 53 recursive approach. This involves the operation of multiple
 54 MTA instances deployed at different levels and requires the
 55 mediation of XCI components deployed over the VI itself
 56 (further details are provided in Section VI).
 57

58 At a lower level of service coordination, the NFVO in the
 59 XCI is responsible for the instantiation of the different NS

components, based on the descriptors and metadata provided
 at the instantiation stage by the tenant. The NFVO, with
 the optional cooperation of the MTA, takes decisions about
 the most convenient usage of infrastructure resources and
 allocates the required VMs and network connections accord-
 ingly. Moreover, during the NS lifecycle, the NFVO is also
 responsible for the continuous monitoring of resource failures
 or infrastructure and application performance, coordinating
 the automated reactions for up/downscaling and self-healing
 procedures at single VNF and global NS level.

V. REQUIREMENTS AND ENABLING TECHNOLOGIES FOR MULTI-TENANCY

Multi-tenancy support requires a coordinated, holistic ap-
 proach from the hardware to the XCI controllers up to the
 application layer, where the MTA acts as a global orchestrating
 entity. In this section we present the main requirements to sup-
 port multi-tenancy at all these layers, analysing the approaches
 that can be adopted to meet them.

A. Data plane

When carrying the data of several tenants through the
 network, several requirements have to be considered:

- *Traffic separation.* One tenant should not be able to listen to the traffic of other tenants or of the network provider.
- *Traffic isolation.* The network has to provide guaranteed QoS to traffic of different tenants. Traffic of one tenant should not impact the QoS of the traffic of other tenants.
- *Traffic differentiation.* The traffic of different tenants may be forwarded differently, even when entering or exiting the network at the same points of attachment.
- *Statistical multiplexing.* Multiplexing gains should be possible among the traffic of different tenants.

The technical solution for *traffic separation and isolation* depends on the specific data plane technology adopted for the XFE, circuit or packet switched forwarding. For circuit switched forwarding, traffic separation, isolation and differentiation can be achieved by creating different circuits per tenant. Although this is beneficial to achieve low and deterministic latency for example, it does not provide *statistical multiplexing* gains among the traffic of different tenants, which are instead enabled with packet switching technologies.

For packet switched forwarding, the multi-tenancy requirements are supported by using a common frame format across the network and different transmission technologies: the 5G-Crosshaul XCF. We propose Provider Backbone Bridge-Traffic Engineering (PBB-TE) [9] as common format to encapsulate the tenants traffic, but other frame formats such as Multi-Protocol Label Switching - Transport profile (MPLS-TP) can be used alternatively (for a comparison of PBB-TE and MPLS-TP see [10]).

In our solution, the fields in the PBB-TE header (Fig. 3) are used to achieve the multi-tenancy requirements as follows. *Traffic separation* is based on the Backbone VLAN ID (B-VID) and the Service ID (I-SID), used to identify the traffic for different tenants by using unique identifiers per tenant

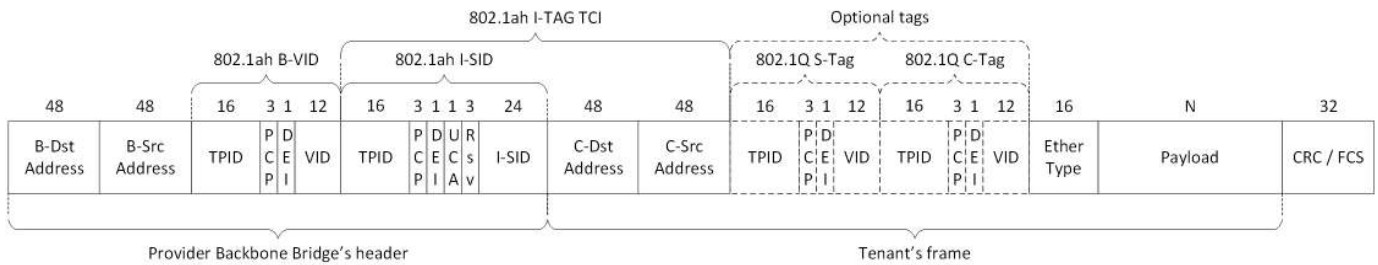


Fig. 3. Provider Backbone Bridge Traffic Engineering (PBB-TE) header

or even per service of the tenants. This allows to create different virtual networks and to keep the traffic separate at the XFEs. Independent forwarding decisions are also taken at the level of these separate traffic flows, thus achieving traffic differentiation on a per-tenant basis. Traffic isolation regarding QoS is based on the three priority-code-point bits within the header, used to distinguish different types of service within the network and to schedule the packets for forwarding based on this priority information. At the ingress of the network this priority has to be set appropriately and consistently across the different tenants to simplify the rules within the network.

Per-tenant XCF forwarding decisions are elaborated at the control plane and configured on the data plane following a forwarding abstraction model common to all the XFEs, either circuit or packet based. Such models are defined by the southbound protocols that define the interaction between the data and control planes. We propose the use of OpenFlow protocols suite as southbound interface for controlling the forwarding of XCF frames.

B. Control plane

Support of multi-tenancy has a strong impact on the XCI components, from the network controller, to the VIM and MANO components for the orchestration and delivery of VNFs and NSs.

At the SDN controller level, multi-tenancy requirements are related to the following aspects:

- *Delivery of per-tenant virtual network infrastructures*, providing the user with a uniform, abstract and data-plane independent view of its own logical elements, while hiding the visibility of other coexisting virtual networks.
- *Logical partitioning of physical resources* to allocate logical and isolated network elements handling per-tenant traffic.
- *Configuration of traffic forwarding* at the data plane level compliant with per-tenant traffic separation, isolation and differentiation in the data plane.

Tenant-based *virtual networks delivery* is handled through a dedicated SDN controller service. Its north-bound APIs allow authorized tenants to request and operate their own network instances following abstract specifications, e.g., based on intent-based network models. Access to virtual resources is wrapped by the SDN controller and it is regulated at the north bound APIs based on tenants profiles. *Physical*

resource partitioning is managed within the SDN controller service through resource allocation algorithms combined with procedures to map logical network concepts with their corresponding entities or traffic configurations at the physical level. *Traffic separation* is achieved through the creation of tagged connections, exploiting the XCF multi-tenant features as explained in section V-A. Forwarding rules for the resulting traffic flows are then installed across the physical network following the paths computed by the resource allocation algorithms on a per-tenant basis (*traffic differentiation*), while QoS is handled through the creation of meters or queues (*traffic isolation*).

An example of SDN application for provisioning of multi-tenant virtual network infrastructures is the OpenDaylight Virtual Tenant Network (VTN) project [11]. The VTN application allows a tenant to request a virtual network. The mapping between network packets exchanged between OpenFlow switches at the data plane and instances of virtual networks defined at the logical level is based on ports and/or VLANs (see Fig. 4). Each virtual network entity implements the typical functions of a corresponding physical element (e.g. virtual routers provide routing, ARP learning and DHCP relay agent functions). Moreover, the tenant has the possibility to control the network behaviour defining a set of actions for flows matching L2-L3 filters.

At the VIM and VNF MANO level, beyond similar considerations on virtual resource allocation and isolation extended to computing elements, suitable modelling of the tenant and its capabilities needs to be supported. *Resource allocation* is handled through the creation of virtual machines and software switches assigned to specific tenants within the XPU, with isolation managed allocating specific addressing spaces and configuring proper routing rules and security groups. *Tenant profiles* are defined at the VIM and at the NFV Orchestrator. At the VIM, each tenant has its own view of the VIM capacity, policies to regulate the access to the resources (e.g. a quota of dedicated resources) and, optionally, custom resource flavors and VM images [12]. Requests for new VI must be authenticated and authorized, and they are evaluated based on the resources still available in the tenants quota. Finally, the access to the instantiated VI is strictly limited to the tenant owing the specific instance. Most of the cloud computing platforms (e.g. VMware, OpenStack) support multi-tenancy.

A similar approach, based on per-tenant profiles and poli-

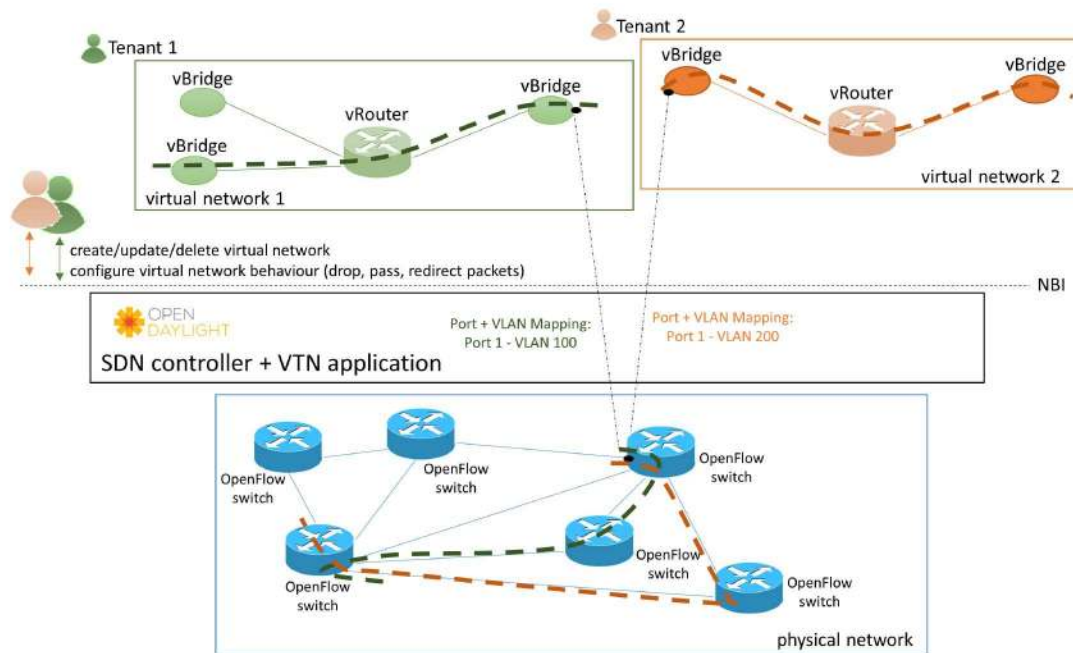


Fig. 4. Virtual networks mapping in OpenDaylight Virtual Tenant Network (VTN) application

cies, needs to be adopted at the NFV Orchestration level, extending the virtual resources concept to VNF and NS entities. Each tenant must have the view and the control on its own VNFs and NSs only; they must be maintained fully isolated from other entities belonging to different tenants, to guarantee their security and their desired KPI level independently on the load of other VNFs. New service requests must be granted depending on the tenants profile, in combination with the tenant-related policies at the VIM level. Currently, this implies extending the functions of the NFVOs such as Open Source Mano⁶ and OpenBaton⁷ to manage tenant separation and mapping between tenants and NSs.

In general, our MTA approach is based on virtualization and this usually involves refinements in the components architecture, enabling one-to-many and many-to-many relationships of software components and implementing the required mechanisms to guarantee security and isolation. From the point of view of performance, the overhead strongly depends on the underlying infrastructure and technology support (VLAN tagging, separate switching instances, compute resource quotas, etc.) and the need or not to emulate such features purely in software. In our considered use cases it is largely within acceptable operational ranges.

C. Application plane

A coherent management of multi-tenancy is required horizontally for unifying the concepts of infrastructure virtualization and multi-tenancy in all involved segments and resources. The MTA at the application level provides such management, becoming the logical decision entity and serving as an optimizer to decide the allocation/modification/deallocation of network, compute and storage resources. Essentially, the application decides an optimum subset of nodes (node mapping) and links (link mapping) in the substrate network to build a VI for a tenant which satisfies its resource demand and SLAs, by solving classical *virtual network embedding (VNE)* problems.

A VNE process consists of two coupled subproblems: node mapping and link mapping problem. The node mapping problem consists of reserving, for each virtual node, enough computational resources of a substrate node without exceeding capacity. Analogously, the link mapping phase consists of finding, for each pair of virtual nodes, a path (a collection of substrate links) to connect them. The selected paths must satisfy the networking requirements of each virtual link without exceeding network capacity on the physical links. The problem is recognized as NP-hard and several approaches (e.g. [13]) have been proposed, which compromise optimality to find feasible solutions.

To deploy and enforce the computed mapping, the MTA needs to interact/coordinate with several functional entities inside the XCI, namely, the SDN controller, the NFVO and the VIM, either to collect information (GET command) or to

⁶<https://osm.etsi.org/>

⁷<http://openbaton.github.io/>

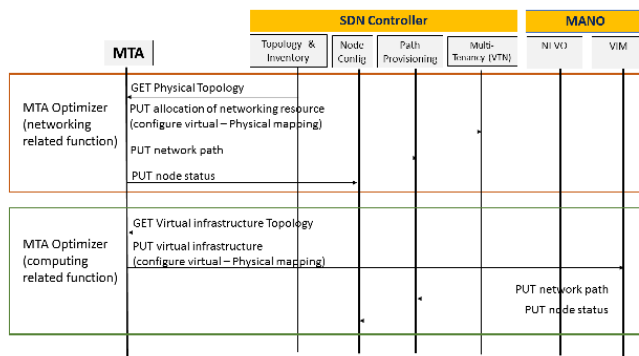


Fig. 5. Workflow of Multi-Tenancy Application: interaction with control plane

provide commands (PUT command) (see Fig. 5). The MTA covers both network- and computing- related functions. The actual workflows are strongly dependent on each use case. For the network-related services, the MTA firstly collects information on physical topology, traffic paths and link load through the XCI, then computes the optimum allocation of networking resources and commands the XCI to perform the required configuration. This may involve direct requests to the SDN controller (to provision network paths and/or to allocate virtual nodes providing the desired mapping between physical and virtual ports). For the computing-related services, the MTA may ask the NFVO to provide a virtual infrastructure topology specifying where the VNFs must be placed or instruct directly the VIM to enforce the mapping between virtual infrastructures and corresponding physical resource. The VIM itself will in turn request the SDN controller for the provisioning of required network paths and related node configurations.

VI. MULTI-TENANCY RECURSION

The 5G-Crosshaul architecture has been designed to not only support sharing of the common transport infrastructure by multiple tenants, but also to allow each tenant to own and deploy its own MANO system. We refer to this case as **Multi-MANO**, building a hierarchy of tenants operating on top of slices of virtual infrastructure. This concept requires support for XCI recursion to allow multiple instances of the 5G-Crosshaul MANO operating on top of the set of services provided by the XCI instance below. The 5G-Crosshaul architecture enables this functionality by, on the one hand, providing support and book-keeping of resources, maintaining a consolidated state of the virtual resources provided to each tenant and, on the other hand, by providing a homogeneous API for controlling the underlying virtual resources which is transparent to the level of the hierarchy where the tenant is operating.

Fig. 6 shows the recursive architecture. In the lower layer, the owner of the physical resources (MNO), instantiates its XCI. Different tenants request the provisioning of virtual infrastructures to the MTA. By means of a template, blueprint or SLA, each tenant specifies not only the slice characteristics

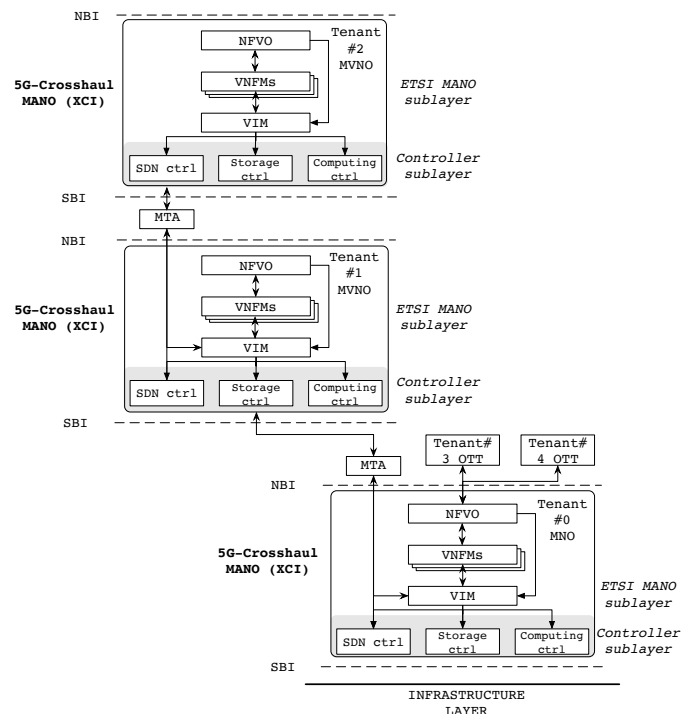


Fig. 6. Crosshaul Control Infrastructure (XCI) Recursion: Multi-MANO

(topology, QoS, etc.) but also some extended attributes such as the level of resiliency desired. The provider must take care of meeting the requirements and managing the available resources. Through the use of the MTA application, the resources at the MNO are hidden to the MVNOs, providing a layer of abstraction easing the management of each slice.

In a recursive and hierarchical manner, each tenant can operate its VI as the MNO operates on the physical one, allocating and reselling part of the resources to other MVNOs. Fig. 6 shows this practice between Tenant#1 and Tenant#2, the infrastructure of MVNO#2 operates over the virtual network offered by the MVNO#1 which operates on top of the MNO infrastructure (the physical one).

The multi-tenant architecture presented in this section is very challenging and one of the central points of innovation of the project. To devise a feasible and flexible framework we have followed the recursion principles of the ONF architecture [14]. Although here we presented mainly the control plane related issues, enabling multi-tenancy in such an architecture also requires modifications in the data plane. For example, isolation of resources/traffic is required, as described in section V-A. In addition, we have designed a specific OpenFlow based pipeline to deal with the forwarding of traffic with different requirements and resources per tenant.

VII. SUMMARY AND CONCLUSIONS

5G requires a new generation multi-tenant transport network integrating fronthaul and backhaul segments into a single transport infrastructure. In this paper we presented the 5G

Transport Network architecture designed in the 5G-Crosshaul project that enables multi-tenancy through network slicing.

We considered two main network slicing services that enable different control and automation levels of network slices management: i) provisioning of Virtual Infrastructures (VI) under the control and operation of different tenants and, ii) provisioning of tenant's owned Network Services (NS) as defined by ETSI NFV. The former deals with the allocation and deallocation of VIs, logical entities encompassing a set of compute and storage resources interconnected by a virtual, logical network. In the latter, NSs are instantiated directly over a shared infrastructure, and as a set of interrelated Virtual Network Functions (VNFs).

A Multi-Tenancy Application (MTA) building on the network slicing services has been described that coordinates and manages the tenants access to the shared infrastructure, performs resource isolation between instances assigned, and delivers related services, such as the allocation and operation of VIs or NSs, by means of a set of proposed APIs.

Finally, the multi-tenancy recursion case (multi-MANO) has been considered which requires to support multiple instances of the 5G-Crosshaul MANO simultaneously.

VIII. ACKNOWLEDGMENT

Work partially funded by the European Commission through the EU H2020 5G-Crosshaul Project (grant no. 671598).

REFERENCES

- [1] S. Gonzalez, A. de la Oliva, X. Costa-Pérez, A. Di Giglio, F. Cavaliere, T. Dei, X. Li, and A. Mourad, "5g-crosshaul: An sdn/nfv control and data plane architecture for the 5g integrated fronthaul/backhaul," *Transactions on Emerging Telecommunications Technologies*, vol. 27, no. 9, pp. 1196–1205, 2016.
- [2] 5G PPP Architecture Working Group, *View on 5G Architecture*, July 2016.
- [3] Next Generation Mobile Networks Alliance (NGMN), "5G White Paper and Description of Network Slicing Concept," 2015/2016.
- [4] K. Samdanis, X. Costa-Pérez, and V. Sciancalepore, "From network sharing to multi-tenancy: The 5g network slice broker," *IEEE Communications Magazine*, vol. 54, no. 7, pp. 32–39, July 2016.
- [5] ETSI, Network Functions Virtualisation, *Network Functions Virtualisation (NFV); Management and Orchestration*, December 2014.
- [6] X. Costa-Pérez, A. Garcia-Saavedra, X. Li, A. de la Oliva, P. Iovanna, T. Deiss, A. di Giglio, and A. Mourad, "5G-Crosshaul: An SDN/NFV Integrated Fronthaul/Backhaul Transport Network Architecture," *IEEE Wireless Communications Magazine*, to be published in Feb. 2017.
- [7] C. J. Bernardos, B. P. Ger, M. Di Girolamo, A. Kern, B. Martini, and I. Vaishnavi, "5GEX: realising a Europe-wide multi-domain framework for software-defined infrastructures," *Transactions on Emerging Telecommunications Technologies*, vol. 27, no. 9, pp. 1271–1280, 2016.
- [8] OASIS, Organization for the Advancement of Structured Information Standards, *TOSCA Simple Profile for Network Functions Virtualization (NFV) Version 1.0*, March 2016.
- [9] IEEE 802.1 Task Group, *IEEE 802.1ah-2008 - IEEE Standard for Local and metropolitan area networks – Virtual Bridged Local Area Networks Amendment 7: Provider Backbone Bridges*.
- [10] R. Vaishampayan, A. Gumaste, S. Rana, and N. Ghani, "Application driven comparison of t-mpls/mps-tp and pbb-te - driver choices for carrier ethernet," in *IEEE INFOCOM Workshops 2009*, April 2009.
- [11] "OpenDaylight Virtual Tenant Network (VNT)," Access on 2017-Feb-10. [Online]. Available: <https://wiki.opendaylight.org/view/VTN:Main>
- [12] R. Vilalta, A. Mayoral, R. Casellas, R. Martinez, and R. Munoz, "Experimental Demonstration of Distributed Multi-tenant Cloud/Fog and Heterogeneous SDN/NFV Orchestration for 5G Services," *European Conference on Networks and Communications (EUCNc)*, 2016.
- [13] L. Gong, Y. Wen, Z. Zhu, and T. Lee, "Toward profit-seeking virtual network embedding algorithm via global resource capacity," in *INFOCOM, 2014 Proceedings IEEE*, April 2014, pp. 1–9.
- [14] Open Networking Foundation (ONF), "SDN Architecture, Issue 1.1," 2016.

BIOGRAPHIES

Xi Li received her M.Sc in 2002 at TU Dresden and Ph.D. in 2009 at University of Bremen, Germany. From 2003 to 2014 she worked as a research fellow and lecturer in Communication Networks Group at University of Bremen, leading several industrial and European research projects on mobile networks. In 2014 she worked as a solution designer in Telefonica, Germany. Since 2015 she joined NEC Laboratories Europe, working as a senior researcher on 5G Networks.

Ramon Casellas graduated in telecommunications in 1999 by UPC and ENST. He worked as an undergraduate researcher at France Telecom and BT Labs, completed a Ph.D. in 2002 and worked as an associate professor (ENST) until joining CTTC in 2006. He is a senior researcher, involved in research and technology transfer projects. His research interests include network control, GMPLS/PCE and SDN/NFV. He has co-authored over 150 papers, 4 book chapters and 4 IETF RFCs.

Giada Landi holds an Italian Laurea degree "cum laude" in Telecommunication Engineering from the University of Pisa. She holds 10+ years experience in telecommunication networks, with focus on control plane architectures and protocols. Currently, she is R&D project manager at Nextworks. Her main research areas include SDN/NFV, cloud computing and orchestration, ASON/GMPLS and PCE architectures. She participated in industrial projects and consultancies on PCE, SDN and NFV topics and is active in European research projects.

Antonio de la Oliva obtained the Degree on Telecommunication Engineering by UC3M in December 2004 and a Ph.D. on Telematics Engineering in July 2008. He has served as Vice-chair of the IEEE 802.21b task group and Technical Editor of IEEE 802.21d, and now he is serving as Rapporteur of the ETSI WP on Crosshauling. He is currently working as deputy project manager of the H2020 5GPPP 5G-Crosshaul project.

Xavier Costa-Pérez is head of 5G Networks R&D at NEC Laboratories Europe, where he manages several projects focused on 5G mobile core, backhaul/fronthaul, and access networks. He is a 5GPPP Technology Board member and Technical Manager of the 5G-Crosshaul project. His team contributes to projects for NEC products roadmap evolution, to European Commission research collaborative projects and related standardization bodies. He received his M.Sc. and Ph.D. in telecommunications from the Polytechnic University of Catalonia (UPC-BarcelonaTech).

Andres Garcia-Saavedra received his M.Sc and Ph.D. from University Carlos III of Madrid in 2010 and 2013, respectively. He then joined the Hamilton Institute in Trinity College Dublin, Ireland, as a Research Fellow in 2014. Since July 2015 he is working as a Research Scientist at NEC Laboratories Europe. His research interests lie in the

1 application of fundamental mathematics to real-life computer
2 communications systems and the design and prototyping of
3 wireless communication systems and protocols.
4

5 **Thomas Deiß** received his degree in computer science in
6 1990 and his Ph.D. in 1999 from the University of Kaiser-
7 slautern. He joined Nokia in 1999. He has contributed to
8 standardization in the area of automated testing and worked
9 in requirements engineering for backhaul functionality of
10 WCDMA and LTE base stations with a focus on backhaul
11 sharing among radio technologies.

12 **Luca Cominardi** received his B.S. and M.Sc. degrees in
13 Computer Science from University of Brescia, Italy. He also
14 received a M.Sc. degree in Telematics Engineering from the
15 University Carlos III, Madrid, where he is currently pursuing
16 a Ph.D. in the same field. He has six years of experience
17 in European research projects, such as FP7 FLAVIA, FP7
18 ICT iJOIN, and H2020 5G-Crosshaul, and he is now with
19 InterDigital Europe as 5G Radio Access Network researcher.

20 **Ricard Vilalta** has a telecommunications engineering de-
21 gree (2007) and Ph.D. degree (2013), at UPC, Spain. Since
22 2010, he is researcher at CTTC, in the Communication Net-
23 works Division. He is currently a Research Associate at Open
24 Networking Foundation. His research is focused on SDN/NFV,
25 Network Virtualization and Network Orchestration. He has
26 been involved in international, EU, national and industrial
27 research projects, and published more than 100 journals,
28 conference papers and invited talks.
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60