Grand Agreement No: 101016941

Full Title: Open cooperative 5G experimentation platforms for the industrial sector NetApps

Start date: 01/01/2021

End date: 31/12/2023

Duration: 36 Months

Project URL: https://www.5g-induce.eu/

Deliverable D2.1

5G platform design and requirements in support of Industrial sector NetApps

Document type	Deliverable		
Title	D2.1 – 5G platform design and requirements in support of Industrial sector NetApps		
Contractual due date	31/12/2021 (M12)	**Actual submission date**	25/01/2022 (v1.0) 04/04/2022 (v2.0)
Nature	Report	**Dissemination Level**	Public
Lead Beneficiary	UBITECH		
Responsible Author	D. Klonidis (UBITECH)		
Contributions from	D. Klonidis, T. Xirofotos, P. Gouvas (UBITECH), C. Lessi (OTE), C. Lombardo, R. Bruschi (CNIT), M. Giribaldi (INFO), Qi Wang, J. Alcaraz-Calero (UWS), S. Mazarakis, E. Petrakou (UNIS), M. Molla (ERC)		

Revision history

Version	Issue Date	Changes	Contributor(s)
v0.1	02/11/2021	Detailed structure and assignment of sections.	D. Klonidis (UBITECH)
v0.2	03/12/2021	Editing of sections 1, 2.2, 2.3.	D. Klonidis, T. Xirofotos (UBITECH), C. Lombardo (CNIT)
v0.3	06/12/2021	Editing of 2.1 on NetApp definition	Qi Wang, J. Alcaraz-Calero (UWS)
v0.4	10/12/2021	Editing of 4.1. Copy of requirements from MS1 and update in 4.2.	D. Klonidis (UBITECH) + MS1 control point: CNIT, INFO, WHR, ERC, 8BELLS, K3Y
v0.4a	17/12/2021	Comments on requirements list. Section 4.3.	C. Lessi (OTE), D. Klonidis (UBITECH)
v0.5	29/12/2021	Editing of 4.4 Non-functional requirements.	S. Mazarakis, E. Petrakou (UNIS)
v0.6	07/01/2022	Editing of section 5 (Business requirements).	D. Klonidis (UBITECH), C. Lessi (OTE), M. Giribaldi (INFO)
v0.7	14/01/2022	Editing of Section 3 (Reference architecture).	P. Gouvas (UBITECH), R. Bruschi (CNIT)
v0.8	21/01/2022	Final editing and revisions. Executive summary and conclusions. Version sent for interval review.	D. Klonidis (UBITECH)
v0.8a	24/01/2022	Revisions and comments.	F. Davoli (CNIT), C. Lessi (OTE)
v0.9	24/01/2022	Version including revisions and comment.	D. Klonidis (UBITECH)
v1.0	**25/01/2022**	Final review and formatting.	F. Davoli, R. Rapuzzi (CNIT)
v1.1	10/03/2022	New additions: Section 2.2 presenting other NetApp definition efforts. Section 2.5 providing the state-of-the-art with respect to service level (NAO and OSS) parts.	D. Klonidis (UBITECH)
v1.2	11/03/2022	Reordering of requirements to start with business requirements first. Changes in the section 4 intro, executive summary, and conclusion to match the new additions and the reordering of sub-sections. Addition of the security requirements in the non-functional requirements.	D. Klonidis (UBITECH)
v1.3	14/03/2022	Revisions and comments.	F. Davoli (CNIT), C. Lessi (OTE)
v2.0	**04/04/2022**	Formatting revised version for resubmission.	F. Davoli, R. Rapuzzi (CNIT)
v2.1	02/05/2022	Added missing contribution information from ERC.	M. Lorenzo (ERC)

Table of Contents

List of Figures

List of Tables

Glossary of terms and abbreviations used

Abbreviation / Term	Description
AGV	Automated Guided Vehicle
AI	Artificial Intelligence
API	Application Programmable Interface
AR	Augmented Reality
BSR	Business Specific Requirements
CAPIF	Common API Framework
CPU	Central Processing Unit
CS	Cloud Service
CVM	Common Validation Metric
DN	Data Network
DRAM	Dynamic Random Access Memory
E2E	End-to-End
ETSI	European Telecommunications Standards Institute
ExFa	Experimentation Facility
GPR	Generic Platform Requirements
GPU	Graphics Processing Unit
GSMA	Global System for Mobile Communications Association
GST	General Slice Template
GUI	Graphical User Interface
HTTPS	Hypertext Transfer Protocol Secure
IaaS	Infrastructure-as-a-Service
ISO/IEC	International Standardization Organization / International Electrotechnical Commission
JSON	JavaScript Object Notation
KPI	Key Performance Indicator
MaaS	Metal-as-a-Service
MANO	Management and Orchestration
MEC	Multi-Access Edge Computing
ML	Machine Learning
MNO	Mobile Network Operator
MSR	Module Specific Requirements
NAO	Network Application Orchestrator
NAT	Network Address Translation
NEF	Network Exposure Function

Abbreviation / Term	Description
NetApp	Network Application
NEST	Network Slice Template
NIC	Network Interface Controller
ODA	Open Digital Architecture
OSM	Open Source MANO
OSS	Operation Support System
PaaS	Platform-as-a-Service
PLC	Programmable Logic Controller
PNF	Physical Network Function
QoE	Quality of Experience
QoS	Quality of Service
RAM	Random Access Memory
RAN	Radio Access Network
RSSI	Received Signal Strength Indicator
SLA	Service Level Agreement
SLAM	Simultaneous Localization and Mapping
STUN	Session Traversal Utilities for NAT
TCP	Transmission (Transport) Control Protocol
TMF	Tele Management Forum
TRL	Technology Readiness Level
TURN	Traversal Using Relays around NAT
UAV	Unmanned Aerial Vehicle
UDP	User Datagram Protocol
UE	User End
VAF	Virtual Application function
VNF	Virtual Network Function
VR	Virtual Reality
WSS	Web Socket Server

Executive Summary

The core outcomes of the design phase of the 5G-INDUCE project are reported and include the defined requirements that affect the overall implementation of the proposed networked application (NetApp) orchestration platform, and the extracted reference architecture.

The NetApp term is defined as the set of networked virtualizable functions being either application-driven functions (i.e., software components) or network service functions (in support of an application use case), that are deployed and linked together over the infrastructure, according to specific end user resource requests and deployment policies. The NetApps are managed by the 5G-INDUCE orchestration platform through the NetApp Orchestrator and the Operation Support System (OSS) modules. The NAO is responsible for the deployment and real-time management of NetApps.

The overall architecture from the end user layer to the network orchestration and infrastructure level includes:

- The NAO graphical user interface (GUI), which provides the enabling framework for the end users (here, the vertical industry experts and the application developers) for the porting and management of their application components and the features of the applications, without interfering with the network level functions and processes.
- The core part of the NAO, which includes the deployment and real time processing functionalities on the application layer, providing for the first case the translation of end user requests to linked application functions through application graphs, and in the latter case the reconfiguration requests through the analysis of the network and resource monitoring changes and the defined policies.
- The NAO to OSS interface, which allows the bidirectional interaction with the slice-enabled OSS, by communicating the application graph annotated with QoS and operational requirements and receiving the materialised slice for the NetApp, as well as potential reconfiguration requests.
- The core part of the OSS, which is in charge of managing all functions and operations required for the placement of the requested application graphs over a slice, as well as of maintaining the information on all the deployed applications, network services, and available resources.
- The OSS southbound interface to the network orchestrator and the infrastructure, which communicates with an external NFV Orchestrator, defining the needed number of virtual links and resources to be applied for the requested NetApp deployments, and also producing the configuration files and commands for each of the deployed VNFs through the VNF Managers at the NFVO level.

The platform adopts a modular implementation approach that decouples the application layer management procedures from the network layer management, providing application awareness to the underlay slice creation and management and compatibility with any structured slice management and network orchestration solution.

The requirements for the implementation activities of the 5G-INDUCE platform are split into functional, non-functional and business/market related. The functional requirements provide the definitions of the expected platform functionalities (generic functional requirements) and the specific implementation characteristics and functions for each module (module specific requirements). The non-functional requirements follow the ISO/IEC 25010 model and are used to determine which quality characteristics should be considered when evaluating the properties of the 5G-INDUCE platform solution. A mapping between the functional requirements per submodule and the non-functional requirements highlight which of these characteristics are of key importance for moving towards marketable solutions. Finally, the business-related requirements define the main needs of each stakeholder for the development of the targeted solutions. The key stakeholders identified are the application software developers, the network operators and the vertical end users, all posing different types of requirements from their own point of view. The fourth type of a system integrator may also play an important role for joining the initial gap between end users (customers) and infrastructure owners (operators).

REMARKS:

– 1 | The use case requirements are studied in D2.2 and focus on the functional implementation requirements stemming from the specific needs of the various use cases. The targeted evaluation metrics are also defined there. The current deliverable D2.1 and D2.2 complete the design phase of the project.

– 2 | The initial implementation of the reference 5G-INDUCE platform architecture is described in D3.4, which presents mainly the current version of the platform with its enhancements with respect to the 5GPPP MATILDA project architecture. The architecture presented here refers to the overall outcome plan for 5G-INDUCE and the implementations that are realised in the next two releases.

– 3 | The main stakeholders are identified in this deliverable (section 4.5) together with their main roles in the project and specifically the exploitation of the 5G-INDUCE platform. Further details and market related analysis will be provided in the following months through deliverables D2.3 (M18), D7.6 (M24) and D7.7 (M36).

1 Introduction

This section introduces the studies presented in this deliverable and their main goal with respect to the overall project. The adopted approach is explained next and is linked to the project's design phase actions. The relation between the current work and the development activities in 5G-INDUCE follows. Finally, the structure of the deliverable is presented.

1.1 Deliverable Purpose

This deliverable reports on the findings of the work performed under 5G-INDUCE Task 2.1, which initiated the project's activities. Its main purpose is to collect the set of technical requirements that guide the definition and development of the 5G-INDUCE end-to-end orchestration platform for 5G-ready NetApps. The set of requirements regards application and network level orchestration mechanisms, dynamic and isolated network slices creation and management, support of the set of network functions through an OSS system, along with the need for development of relevant northbound APIs, as well as provision of independent service guarantees per NetApp development and deployment use case. The requirements define all the key functionalities of the 5G-INDUCE platform that need to be supported for the realisation of the 5G trials' demonstrators leading to technological and business validation of the introduced industrial use case solutions.

Based on these requirements and the overall vision of the 5G-INDUCE platform architecture, as described in the Grant Agreement Annex 1 part B, the secondary purpose of D2.1 is the extractions of the exact and detailed system design that will be next developed under WP3 and WP4. This design concerns primarily the 5G-INDUCE NAO and OSS part of the platform, the interfacing with the network orchestrator part and the onboarding mechanism for the proper deployment of the NetApp components.

The following list summarizes the main targeted outcomes of the current deliverable D2.1:

- Definition of the key concepts and functional layers for creating a common understanding.
  - NetAPP versus network functions, services and applications.
  - Application orchestration – Operations support system – Slice manager – network orchestrator – Infrastructure.
- Extraction of requirements that define in detail the functionalities of the 5G-INDUCE platform, including:
  - Definition of the user interface and related processes.
  - Definition of the NAO-OSS interface to any standardised underlay network orchestration platform.
  - Definition of the interfaces between the platform modules contributing to a clear flow of information.
- Extraction of detailed requirements per module that will assist in the development of the specific solutions.
- Design and description of the overall system reference architecture.
  - Detailed flow of information diagram for the main processes (application registration, onboarding, slice creation and management, application management, monitoring and reporting).

1.2 Relation with other deliverables and tasks

Deliverable D2.1 summarizes the key outcomes of the 5G-INDUCE design phase concerning the reference platform architecture and the implementation requirements. These outcomes are essential for the development phase of the project and in particular the elements of the core 5G-INDUCE platform implemented in WP3 and WP4. They also define the overall integration process, by highlighting the interfaces among modules as well as the targeted implementation goals. More specifically, D2.1 feeds the following activities:

a. The NAO (T3.1) and OSS (T3.2) modules' implementation, as well as the complete integrated 5G-INDUCE platform (T3.3)
b. The targeted use case applications' conversion into NetApps including the onboarding process (T4.1 and T4.2)
c. The implementation targets for the integration process in T5.3
d. The parallel impact evaluation and market analysis work starting with the identification of the strategic investment sectors and involved stakeholders (T2.3), followed by the NetApp related market analysis studies and later the identification of the business models and evolved market potentials in T7.5.

D2.1 is also linked directly with D2.2, which focuses more on the specific use case requirements and needs considering the 5G-INDUCE platform capabilities and functions.

1.3 Approach

The adopted approach during the design phase of the project (M01-M12) considered 6 main actions that are listed below and also depicted in Figure 1.1:

Action 1. First a clear definition is attempted with respect to the use of the NetApp term and in relationship to: a) the network functions (seen as any virtualizable function at the network operation level) and b) the applications (seen as a set of linked software components and attached services). This assists also in identifying which parts of the platform deals whit what type of information and from which end user type. In turn, a common understanding is built with respect to the targeted reference layers (NAO, OSS and Network orchestration, Infrastructure) and the more specific fields that are handled by the 5G-INDUCE platform.

Action 2. In parallel to the first action, the initial generic architecture is defined in detail with respect to the targeted features identified for the 5G-INDUCE platform and in support of the application use cases. This is required in order to guide the extraction of requirements having an overview of the targeted system functionalities in mind.

Action 3. The next step in the approach takes into consideration the targeted use cases and the diverse set of application requirements that dictates a set of deployment needs and interdependences from edge to core cloud. In this step the consortium starts considering also the potential DevOps testbed and ExFas' capabilities, performing an initial integration planning to identify any related implementation issues that may arise, well ahead of the planned activities in WP5.

Action 4. The extraction of requirements is initiated right after the previous actions and split into generic platform requirements, and module/engine specific requirements. The requirements are mainly driven by the overall project objectives and the initial architecture, while also considering the targeted implementation through the use cases and the optimum interfacing among the modules for the integration stage. The initial requirements update the architecture in action 5. In turn, according to the updated architecture, the final requirements are extracted. The final requirements consider also business and market-related aspects from the potential use of the 5G-INDUCE platform in real system deployments, driven by the use cases. The initially extracted requirements are put together and reported according to MS1.

Action 5. Based on the initial requirements (MS1), the system architecture is defined in detail including the interfacing among the various modules. Important additions to the architecture are also defined at this stage, concluding in a complete flow of information from the user interface where applications are onboarded to the network configuration through the selection of the required slice. The lifecycle management processes are also defined here for a dynamic application driven operation. Moving towards a more detailed architecture, this action provides feedback to action 4 by defining also in more detail the interfacing related functional requirements.

Action 6. The final step is focused on the definition of the use case scenarios according to the identified use cases and with a view towards the overall system architecture and targeted modules. The use case requirements are reported together with the platform functional requirements under MS1. Then, they are split from the platform requirements and feed the NetApp development activities in WP4 while being reported separately in D2.2, in parallel to the current one.

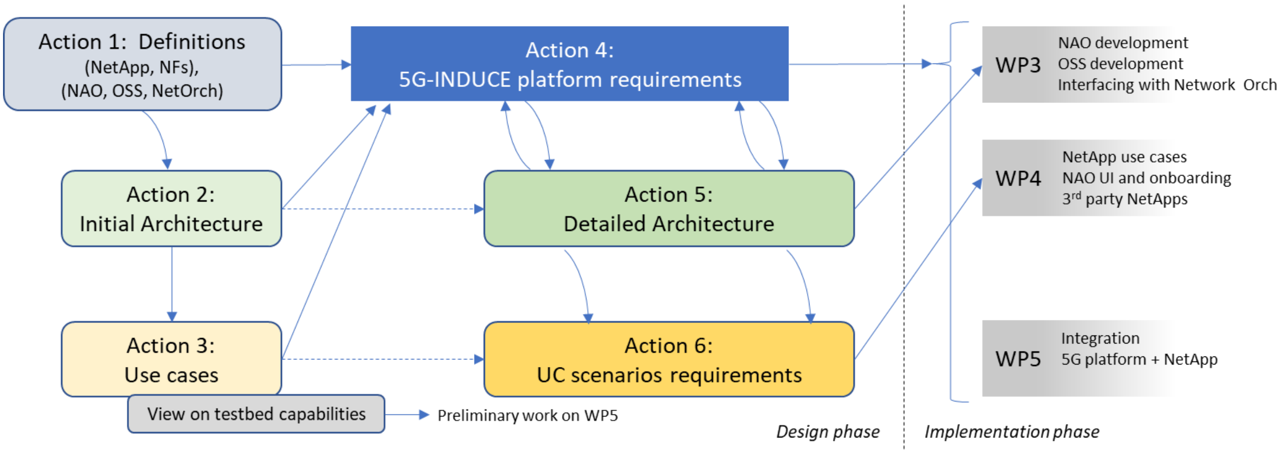

Figure 1.1: Overall approach for the design phase of the 5G-INDUCE project

1.4 Structure of the deliverable

The deliverable is structured according to the key outcomes of the studies. It is noted that for a clearer presentation, the extracted reference architecture section precedes the requirements section. This allows the reader to gain a better understanding of the overall architecture before focusing on the detailed implementation requirements and the specificities of each targeted module.

Following the current introductory section 1, the deliverable is structured as follows:

- Section 2 provides the definition of the NetApp with respect to network functions, as well as the definition of the 5G-INDUCE platform working layers and interfacing layers.
- Section 3 presents the detailed 5G-INDUCE reference architecture, including the platform's modules, the interfaces with the end users and the network orchestrator, and the flow of information.
- Section 4 focuses on the requirements, starting with the overall approach and syntax and followed by the elaborated list of functional and non-functional requirements. A summary commenting on the key outcomes is included also at the end of this section.
- Finally, section 5 concludes on the main findings of the proposal.

2 Definitions and state of the art

The term NetApp is introduced in all 5GPPP projects funded under the H2020 ICT-41 call, but at the time of issuing of the current deliverable there is no detailed definition of this term in the standards and the broader 5G community. For this reason, it is important to present first the 5G-INDUCE view on the definition of a NetApp and in conjunction with the known term of VNF and under the generic scope of a 5G application. Other related NetApp definition activities are presented and linked to the 5G-INDUCE approach. In addition, this section defines the key features of the NetApp Orchestrator (NAO) and the related Operation Support System (OSS) that comprise the main elements of the 5G-INDUCE platform for the onboarding, deployment, management and configuration of NetApps over 5G enabled infrastructures. Finally, an overview of the current state of the art related to the targeted 5G-INDUCE platform topics is provided.

2.1 NetApp definition and VNF types

In 5G-INDUCE, a NetApp is defined as follows:

NetApp is a set of networked virtualizable functions, together with the required resources (compute, storage, network etc.), deployable and operating over 5G and beyond network infrastructures, and being distributed across the network continuum.

The **virtualizable functions** in a 5G system are split in the following three types:

1. The Customer-facing service virtualizable functions, denoted as virtual application functions VAFs – referring to the application functions that are typically developed by vertical service developers for particular vertical use cases (see section 4.5) and linked together in order to provide the required service functionality.
2. The Network service and resource-facing virtualizable functions, denoted commonly in standards as virtual network functions VNFs – referring to the standardised functions provided by the network operator and infrastructure owner for fulfilling the networking and connectivity services (e.g. User Plane Function (UPF), Access and Mobility Function (AMF), Session Management Functions (SMF), etc.).
3. The Value-added (middleware) VNFs or non-standardised VNFs – referring to network-level functions that are added in support of the networking operations and services either to satisfy specific end user network function requests (i.e., intermediately on demand) or as a response to network changes (i.e., automatically). examples may include encryption and decryption VNFs for secure communications, video processing VNFs and so on.

The considered distributed **network infrastructure** spans from Edge, Transport network, Core network and Data Networks (DNs) to Cloud networks.

➔ A NetApp SHALL employ VAFs, and MAY additionally utilise value-added VNFs.

➔ A NetApp SHALL be deployable over any part of the network infrastructure continuum

REMARK – A: Use case applications are not NetApps. They become a NetApp once they are deployed over a network infrastructure and include specific networking and policy requirements. Therefore, NetApps are linked to applications but refer to the deployed versions of applications that are further extended with the required resources and VNFs over the targeted infrastructures.

REMARK – B: The applications are composed by one or typically many software components. These applications components in containerised deployments are becoming the VAFs. Therefore, the VAFs have practically the same notion as the application components. Once the NetApp is deployed, the linked and deployed VAFs are becoming the NetApp components.

For further enhancing the definition of NetApps, the following requirements are identified for the creation of NetApps:

- A NetApp is deployed on demand as requested by a vertical end user. A service provider may manage the deployment on behalf of a vertical end user customer.
- An infrastructure owner (typical a network operator) deploys the requested NetApp over the infrastructure according to the NetApp deployment requirements and the polices.
- A NetApp must be cloud-native compliant, to allow automated cloud-based deployment and benefit from the cloud deployment such as built-in elasticity.
- A NetApp may operate over a network slice (over a single or multiple administrative domains).
- A NetApp may have embedded intelligence in its components (i.e., virtualizable functions) and/or benefit from intelligence provided by Multi-access/Mobile Edge Computing (MEC) and/or Core network, etc.
- A NetApp must be compatible with various networks such as private 5G and beyond networks, and hybrid private and public 5G and beyond networks.
- A NetApp may have resource and network requirements in terms of hardware, such as memory, CPU or GPU usage, H/A availability, etc.
- The composition and deployment of a NetApp should be compatible with the various relevant standard architectures, interfaces and APIs, such as ETSI NFV-MANO, 3GPP Service Based Architecture and related network functions for network exposure, network slice selection and so on.
- For security considerations, general network security requirements and mechanisms in future networks SHOULD be applied, and NetApp specific security network functions MAY be included as part of a NetApp.

2.2 Related NetApp definition efforts

As mentioned at the introductory paragraph of this section, there is no detailed definition of the NetApp term in the standards and the broader 5G-community. The term has been initially introduced in Horizon 2020 ICT-41 calls and was referring in general to the management of the service level solutions in the vertical sector.

The initial requirement for a clear definition of the NetApp term has been identified by the 5GPPP Technical Board members. The definition efforts have been passed to the 5GPPP Software Networks WG and initiated as a consolidated version among all ICT-41 projects.

In parallel to the above-mentioned efforts, 5G-INDUCE has worked in close alignment with the ICT-41 project 5GASP (www.5gasp.eu) on proposing a common definition. Later, collaborations have been established also with the ICT-41 projects 5G-IANA (www.5g-iana.eu) and VITAL-5G (www.vital5g.eu).

In 5GASP, the NetApp, in the context of the 5G System, is defined as "set of services that provide certain functionalities to the verticals and their associated use cases" [1]. Similarly to 5G-INDUCE, the set of services is seen as a collection of application components that are blended together with network functions. The NetApp is the deployed set of services (i.e., functions) which can then be tailored to different end-user needs. An important addition that is brought in from 5GASP is the mapping mechanism that defines a certain NetApp template architecture for following newly standardised general slice templates (GSTs) from GSMA and TM Forum. In turn, 5GASP identifies the 5G-INDUCE NAO management deployment mechanism as a key feature for providing a wholistic interface with the end user's vertical services.

In VITAL-5G, the NetApp definition is linked to any type of application- or network-related function that is accordingly parameterised by the end user (i.e., vertical use cases). The adopted concept is essentially similar to 5G-INDUCE and 5GASP, but the granularity level that is adopted is finer and goes into the component level. A second level is also defined with chained components serving a certain end-user service, which is fully aligned to the definition in 5G-INDUCE. The rationale in VITAL-5G is that any software component or chain of components linked to a user-defined vertical service forms a NetApp, once the networking parameters are defined and can be reused in any other service.

In 5G-IANA a mixture of the approaches defined in VITAL-5G and 5G-INDUCE is adopted leveraging the use of a well-defined component and NetApp cataloguing module. Here, once the application or network-oriented components are onboarded followed by the required network parameters, they update the catalogue and can be recalled by any service or the slice management module on demand. The NetApp is though defined for the chained only components that serve an end-user service, and may even include single component NetApps when these can be deployed as simple services (e.g., a server).

2.3 NAO module purpose

The scope of the NAO is to undertake the deployment and real-time management of NetApps, while inherently providing elasticity and compliance with certain high-level NetApp policies. Its purpose is to provide an interfacing layer to the end user (i.e., the vertical industry expert and the application developer) for managing the deployable applications and their features, without interfering with the network level functions and processes managed by the involved telecom operators and infrastructure owners. Thus, NAO decouples the application layer management procedures from the network layer management, providing application awareness to the underlay slice creation and management and compatibility with any structured slice management and network orchestration solution.

The overall concept is aligned to the way that modern complex vertical applications are designed over distributed architectures with edge processing capabilities. Such applications consist of a chain of cloud-native components that can be managed independently, as far as their scaling aspect is concerned. Each application component is bundled in an orchestration-friendly way, i.e., as a VM image, a container, or even a unikernel maintaining a backward compatibility with all three industry-leading approaches while offering important telco-interplay capabilities such as Bi-directional interaction with a slice-enabled OSS, Resource-constrained slices, Application profiling and Policy enforcement.

Having this concept in mind, the NAO is designed to leverage the powerful advancements of the 5G-INDUCE OSS, such as the refined SB convergence layer and the advanced functions on 5G Network Exposure, Network Slice Selection, and Policy Control functions, exposed by the "Network Service Manager". The goal is:

- (i) to realize the interfacing towards a substantially enhanced slicing model (slice intent), suitable for accommodating advanced slicing functionalities for NetApps
- (ii) to virtualize and elegantly expose these resources to the overlay NetApp components, forming a rich NetApp composition API, without disclosing sensitive information (e.g., UE or topology formation data) about the underlying network
- (iii) to support run-time management of the NetApp lifecycle and graph formation in order to dynamically enable, scale and modify NetApp capabilities according to relevant events e.g., the number and type of connected devices/things, their bandwidth on the radio link, their positioning, etc.

2.4 OSS module purpose

The 5G-INDUCE OSS is in charge of managing all functions and operations required for the placement of a NetApp over a slice, as well as maintaining the information on all the deployed applications, network services, and available resources. Along with the NFVO, it forms the telecom layer of the NetApp Management Platform.

The OSS extends the legacy ETSI-MANO specification by exposing, at the north-bound, the network and computing resources required to compose a slice and, at the south-bound, by strongly leveraging on the NFVO to have full control of the lifecycle management of network services and functions.

As will be described in the next section, the OSS is composed of state-of-the-art, parallelizable cloud native software to facilitate integration within service-meshes. The current version builds upon the prototype developed within the MATILDA Project. At this stage of the 5G-INDUCE lifetime, several extensions have already been introduced with the aim of supporting the modification and reconfiguration of the slice within its lifecycle and enabling advanced operations to deal with UE mobility and dynamic QoS/operation NetApp requirements. 5G-INDUCE NetApps will be lively scaled and relocated onto computing facilities in new geographical areas or at different network infrastructure aggregation levels, with their traffic from and to UEs steered accordingly, as well as the platform components and network services. Moreover, the OSS will also manage and terraform bare-metal resources by means of a Metal-as-a-Service (MaaS) approach.

2.5 State of the art with respect to the targeted 5G-INDUCE platform

This sub-section reviews the current state of the art related to the main modules of the 5G-INDUCE architecture as these are defined above. The analysis includes the efforts in the standardization and scientific community with respect to the services and the definition of standardized templates that can be easily adapted to the end-user needs, while providing an efficient mapping to the underlay slice deployment and resource management layers. This is essential, in order to better define the added value features of the 5G-INDUCE architecture (see the next section) and create a broader understanding of the rationale behind the main concept of 5G-INDUCE. At this stage it is important to highlight that the main goal of the 5G-INDUCE platform is to provide the means to the vertical end users (i.e., the industry sector in our case) to onboard, parameterise and deploy their applications following application-specific requirements and certain policy criteria (as defined through SLAs). This denotes the deployment of an application-oriented approach that extends to the deployment and runtime reconfiguration level of virtualized application components over the network infrastructure. Therefore, the main focus of the platform is on the application and service management layer and touches upon the network layer for creating the appropriate interfaces at the infrastructure level.

2.5.1 Standards at the service layer and mapping with slices

With respect to the standardization activities, the focus is primarily on GSMA and TM Forum actions which deal with vertical industry requirements’ mapping into network slicing and the evolving OSS/BSS matters, through the prism of 5G vertical services and needs, and the compliance with standard 3GPP and ETSI network orchestration solutions.

In **GSMA** a main effort is directed towards the creation of a commonly accepted communication model among operators, providers and vertical costumers, with regard to the needs and requirements of vertical use cases. The initial findings of the analysis are summarised in [2] and [3] and provide the first consolidated effort to collect and map such end-user requirements in a structured and expandable manner. The addressed requirements have been quantified and categorized into performance, functional and operational requirements. However, there was no agreement on how verticals would express these requirements in this study. This has been lately addressed with the definition of a generic slice template (GST) model that is introduced in order to describe a network slice, containing all possible attributes required to define a slice independently of the vertical use case type. The GST can be used to create the network slice template (NEST) with vendor-specific values and serve as an agreed reference between the slice customer and operators, while facilitating the definition of network slices across multi-operator roaming agreements. With this approach, the service requirements from specific vertical industry use cases can be translated into technical requirements, following commonly agreed values or value ranges. Once NEST is created and the network requirements are defined, the 3GPP network slice preparation phase can be initiated. Furthermore, it is noted that NESTs are split into Standardized NESTs (S-NESTs) and Private NESTs (P-NESTs). The former provide a mapping to the standardized 3GPP Slice/Service Type (SST) values, introducing the "minimum requirements" to address each Network Slice Type. The latter can be operator-specific NESTs, permitting operators to negotiate with customers the terms of use of the infrastructure.

In **TM Forum** there is a strong focus on evolving OSS/BSS solutions that are driven by verticals and are able to be integrated into existing standards-driven architectural frameworks. TMF introduces the Open Digital Architecture (ODA) [4], which provides scenarios for Business and Infrastructure Functions and their respective implementation through technology neutral "flavours". These implementations are offered in the form of Open APIs, allowing vertical customers to interact and consume offered network slices or network services or even platforms and infrastructures as a Service. The vendor-agnostic nature of Open APIs and their bottom-up composition across layers make TM Forum's Open APIs a perfect candidate for integrating components among multi-vendor environments. This leads currently to a broad adoption from telecommunications industry. From a customer-facing viewpoint, the ODA model makes reference to a set of Open APIs that could be used for slice management, including APIs for service catalogue management [5], service ordering [6], service inventory management [7]. Also, TMF defines the Network as a Service (NaaS) Component Suite Profile [8], which includes the operations required to be exposed in order to provide the functionality required across interworking Operational Domains.

2.5.2 Research activities at the service layer

The efforts from GSMA and TMF provide a structured framework for the mapping of services to network slices which enables the link between services, deployment and life-cycle-management over the infrastructure. However, there is no specific framework for defining the application layer interface and the onboarding of application functions together or separate to the network functions.

The concept of the application orchestrator has been initially introduced in the MATILDA project [9] and linked to a highly modularized OSS layer for managing the different components of the whole 5G infrastructure, through an underlying convergence Layer. Here the application orchestrator had the responsibility to translate the high-level end user requirements into slice intents and in turn to provide the life cycle management of applications based on application layer monitoring information received through the NFV in combination with policy criteria. It is noted that the MALTIDA approach for adopting a separation between the application orchestration layer and the slice management and OSS level is a unique approach within the set of 5GPPP projects and other related activities globally. The typical approach for handling services is to statically introduce them on top of the OSS or more specifically the slice manager layer, allowing in turn any life-cycle management actions (if applicable) to be within the network orchestration layer or the OSS management layer.

With the introduction of the NetApp concept primarily through the 5GPPP ICT-41 projects, see [10], the deployment and management of the use case application functions is highlighted and linked to that of VNFs. While there is no specific standard for describing the NetApp concept, the ETSI NFV model is adopted in various of these project initiatives as the one describing the delivery and deployment of NetApps. The interaction of the NetApps with the underlying 5G fabric is recommended in all cases to follow the 3GPP standards, as for example via 3GPP NEF- Network Exposure Function APIs, consuming them by following the 3GPP CAPIF definition or when exposing APIs to be consumed by other entities. For the mapping of application function requirements to the specificities of the network infrastructure the adoption of the description templates presented above provides a functional framework at least for the projects 5GASP, 5G-IANA and VITAL5G.

Finally, it is worth mentioning the latest research efforts for introducing Zero-Touch Provisioning (ZTP) capabilities in 5G and beyond networks. Today, ZTP is only explored for network devices (i.e., switches, wireless access points, routers, firewalls, etc.), mainly for automated device provisioning and firmware updates. With the advent of software-defined networking (SDN), ZTP drew further attention, as it could be naturally seen as a key procedure of an SDN controller. The latest research efforts foresee the adoption of zero touch provisioning in the real time management of 5G systems, by leveraging on data analytics collected at the management plane and going as far as fully self-driven 5G networks [11]. The goal is to enable network automation, through dynamic, policy-driven network optimization and slice orchestration with minimal input from human operators [12]. In this regard, Machine Learning (ML)-driven zero-touch slice orchestration solutions have been proposed to make adaptive resource slicing decisions with no prior knowledge on the traffic demand. Moreover, cloud-native technologies are considered the future of vertical application development, which has led to an explosive growth of studies on the applications of ML control for the lifecycle management of containerized cloud-native applications [13]. Evolved 5G Networks will have vertical applications deeply ingrained in the 5G ecosystem, disaggregated in micro-services deployable at the 5G network Edge (Multi-access Edge Computing – MEC concept) [14]. However, there is little ongoing work on cross-domain aspects that jointly optimize the 5G network with underlying Slices and MEC applications. The very recent move towards containerised structures allows any NFV virtualization level to be flexibly handled at the OSS layer and NetApps to be deployed with cloud-native containers in the form of PaaS. Monitoring, analytics and automation mechanisms can be therefore linked across the application management and slice management/OSS levels, as well as with end-user driven policies and service level analytics.

3 Reference 5G-INDUCE architecture

This section presents the 5G-INDUCE reference platform architecture and the individual modules and technologies adopted for its realisation. The architecture follows the initial requirements extracted during the design period of the project and presented later in section 3.

3.1 Overview and Walkthrough

5G-INDUCE aims to deliver a holistic platform where 5G NetApps' lifecycle is managed on top of programmable 5G slices. The logical architecture, depicted in Figure 3.1, is separated into two distinct 'swim lanes' or domains, i.e. the ***NetApp Orchestration (hereinafter NAO) Administrative Domain*** and the ***5G Telco Administrative Domain***.

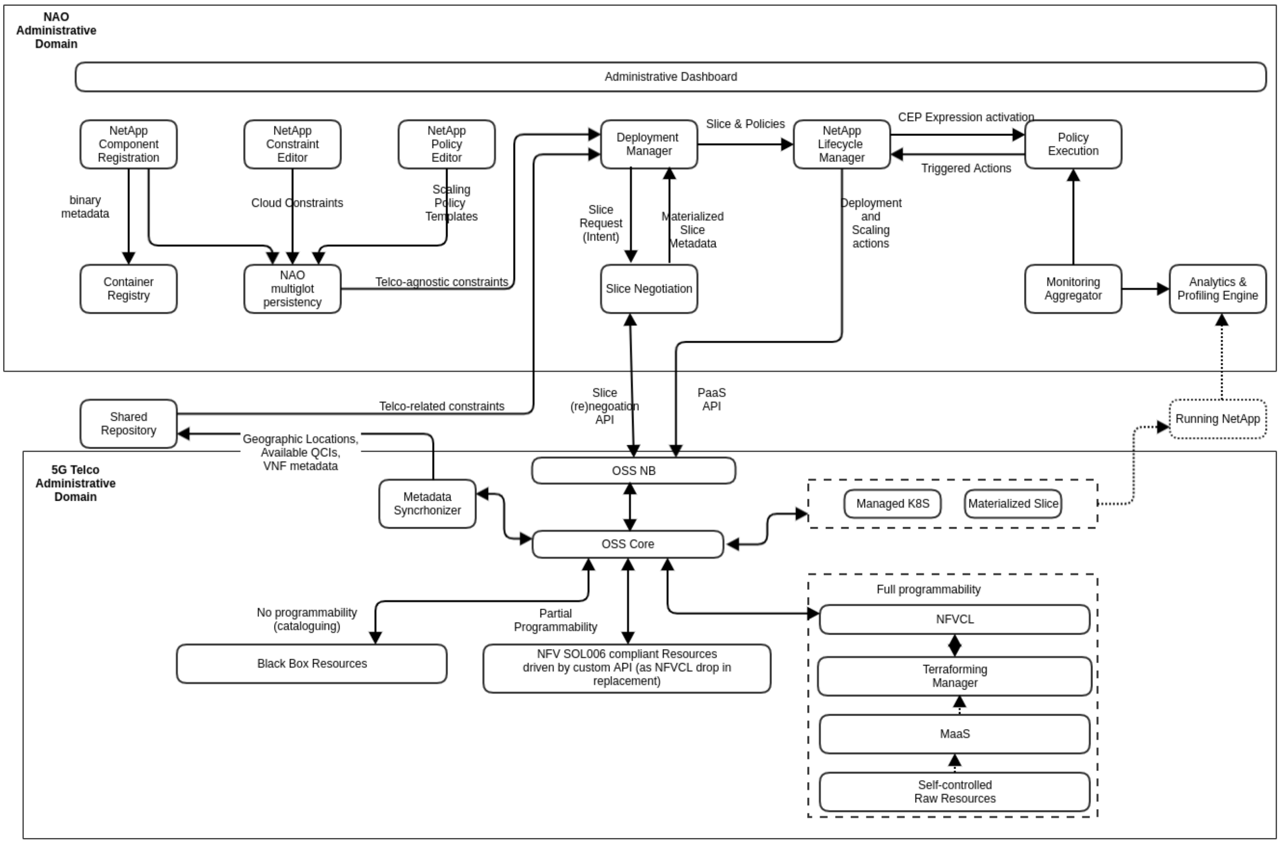

Figure 3.1: Conceptual architecture

The NAO Administrative domain consists of all modules that are responsible for **registering a NetApp and all its components**, **providing cloud-related** and **slice-related** metadata, **authoring deployment and runtime policies**, **negotiating slices** and **managing the operational state** of the NetApp within the scope of a materialized slice. On the other hand, the 5G Administrative domain consists of all modules that are responsible for materializing **a slice based on a request,** thus coordinating all activities that **transform bare-metal and virtual resources** into a programmable environment that can be used by service providers through the NAO.

The walkthrough initiates when a service provider wishes to deploy one of its NetApps' to a 5G environment offered by a telco provider. A NetApp **by definition is a cloud native application that uses several VNFs** as part of its data plane in a seamless way. That is to say that layer-7 business logic is encapsulated in container format; thus, strictly complying to cloud-native principles (a.k.a. 12-factor principles). The NetApp per se has to be registered along a set of strict/formal metadata that dictate the capabilities of the component in terms of scaling and the operational prerequisites that arise during deployment. The service provider is using the ***NetApp Component Registration*** to upload / register and label the NetApp. The registration process is assisted by an intuitive user interface.

Upon registration of the NetApp, its **telco-related requirements** have to be expressed formally. These requirements will be the base for the slice negotiation during a deployment process. These requirements refer to QCI parameters, networking quality thresholds, radio-related aspects and geo-location aspects. Service providers will use the **NetApp Constraint Editor** to author and validate these requirements.

Cloud-related requirements and telco-related requirements drive the slice negotiation process and the initial placement/deployment. However, each NetApp entails a runtime behaviour that affects the way this NetApp scales up and down based on throughput, cloud resource consumption and regional load. The conditions that trigger compensation actions are defined by runtime policies. Such policies are authored using a graphical ***Policy Editor***. The Policy Editor serializes the conditional statements and the defined actions in an abstract format and saves it to a **Polyglot Persistency Engine**. The same storage engine is used to serialize the NetApp cloud and telco-related metadata.

The slice-negotiation process starts when a service provider selects one of the registered NetApps, in order to deploy it in a pre-selected 5G telco provider. Selection of the NetApp implies under the hood the definition of **three distinct 'sets' of information,** i.e. cloud constraints, telco constraints and active policies. After selecting this information, the NAO component initiates a **slice negotiation process** with the selected Telco provider. The negotiation signalling is performed through a Northbound interface of the OSS.

The OSS Core component is responsible to interpret all requirements and perform two major tasks; **a) to infer** whether or not the requirements can be materialized and **b)** to **coordinate underlying resources** in order to expose through the OSS NB API a programmable layer that will be used by the NAO. Programmability implies compute/storage and networking configuration taking under consideration the VNFs that should be deployed. In case slice requirements are materialized in a concrete slice, the NAO undertakes the task to deploy individual components in proper 'placeholders'. The connectivity among the cloud resources will be guaranteed by the **OSS Core** through **underlying programmability that varies according to the given resources**.

The initial deployment of the NetApp results in a given 'state' that can **change based on runtime parameters**. The way that it will change and the triggering conditions are defined by **activated policies**. Such policies define the **conditions** that have to be met in order for specific control plane actions to be performed. Based on activated conditions proper measurements have to be collected by a measurements' collector. The **measurement subsystem, the policy evaluation engine and the NetApp Lifecycle Manager** are continuously operating in the form of a closed loop until the process is deprovisioned.

Indicative control plane actions that can take place include scale in and out (up/down), regional expansion of the service, change of service class, etc. The most complex control-plane action that can occur is slice renegotiation.

3.2 NAO Administrative Domain Components

NetApp Component Registration

The aim of the NetApp Component Registration is to assist the service providers wrap cloud native NetApps in a proper format, so as to be publishable in the Container Registry. Each NetApp consists of multiple containers that are chained in the f**orm of a service graph**. Every cloud-native component has to comply with a specific metamodel which is called the 12-factor metamodel. Compliance provides a proper guarantee that a component will be 'orchestratable' during its deployment. The registration process will provide **design-time validation as far as a component model** is concerned. The registration process shall guarantee that all cloud-native properties are maintained, such as **metadata regarding *minimum infrastructural requirements***, **metadata regarding *deployment preferences***, **metadata regarding *configuration parameters during component initialization***, **mutable configuration parameters during runtime**, **exposed and required interfaces**, **exposed metrics**, ***link metrics*, etc**. The registration process will be performed through a dedicated user interface which will offer several design-time features, in order to make the overall process less error-prone.

NetApp Constraint Editor

As already mentioned, cloud-related constraints are authored through the NetApp Component Registration module. However, there are some **non-cloud-related metadata** that play a significant role in the slice negotiation process. These are 5G-telco related metadata and refer to:

- the link quality between the components of the NetApp, the required radio capabilities (eMBB (enhanced Mobile Broadband), URLLC (Ultra Reliable Low Latency Communications), mMTC (massive Machine Type Communications))
- the QoS Class Identifier classification
- the guaranteed bit rate and
- the required VNFs as part of the VNFFG that is materialized by the telco.

In order to provide these constraints in a formal way, a specialized editor will be developed. It should be noted that all parameters above can be defined without prior interaction with the telco provider, except the required VNFs. During the definition of the required VNFs the tool should know beforehand the exact VNFs that can be supported by the telco. To do so, this component interacts with a **Shared Repository** between the NAO and the Telco Provider. This repository contains metadata of all supported VNFs that can be potentially deployed within a slice; thus, acting as a **shared catalogue**. The responsibility for maintaining this catalogue is delegated to the telco's OSS, as explained later.

NetApp Policy Editor

Every NetApp that is deployed can be subjected to runtime changes/reconfiguration. This reconfiguration aims at the satisfaction of a set of business goals that are bundled in the form of a Service Level Agreements. In the frame of satisfaction, **many types of control plane "actions"** may be required. Indicative actions include allocating more resources, spawning new instances of cloud-native components, migrating live instances, or even **renegotiating the entire slice,** etc.

All these actions will be 'instructed' by proper rules that take under consideration the capabilities of instrumentation (i.e., collection of measurements) and the OSS. These instructions should be provided in a normative way, so as to be executed upon instantiation of a slice. The specific module will be responsible for **authoring these instructions** in a formal rules format. Beyond authoring rules, the Policy Editor will be responsible for ensuring the **structural validity of provided rules and verifying their applicability**. Verification of applicability will be performed by examining the existence of proper "enablers" that will facilitate the execution of a rule.

Polyglot Persistency Engine

This module acts as a persistency layer for: a) the registered NetApps, b) the deployment and slice constraints and c) the authored policies. As a persistency layer, the main functionalities that have to be supported are scalable storage and searchability. Regarding scalability, NetApps along with their metadata have to be stored in a scalable storage engine. Since the execution-ware technology will rely (for most of the cases) on Linux containers, a scalable container storage engine will undertake the task of storing container images.

Regarding searchability, in order for a NetApp to be instantiated it has to be searchable. Searchability will be achieved using a highly responsive scalable indexing engine (e.g., ElasticSearch) which will index the entire serialized component model. Upon indexing, a component can be identified using keywords or metadata.

Slice Negotiation & Deployment Manager

Network slicing is a key feature for the 5G Networking compendium. It is about transforming the network/system from a static "one size fits all" paradigm, to a new paradigm where logical networks/partitions are created, with appropriate isolation, resources and optimized topology to serve a particular purpose or service category or even individual customers. The purpose of the **Slice Negotiation component** is to **handle the complete lifecycle of Network Slice creation**.

The lifecycle of a slice includes **a)** the **slice request** phase, **b)** the **slice instantiation**, **c)** the **slice operation** and **d)** the **slice deprovision**. During the instantiation phase, the creation and verification of the necessary network environment which will be used to support the lifecycle of the NetApp is performed. The environment will be adapted to the **advertised requirements from the Constraint Editor**. The translation of the requirements to a proper slice configuration plan is a primary objective of the **Slice Negotiation**. During the instantiation phase, all resources shared/dedicated to the slice should have been created and configured. The activation step includes any actions that make the slice active, e.g., diverting traffic to it, provisioning compute resources, etc.

NetApp Lifecycle Manager

The NetApp Lifecycle Manager is responsible for maintaining the equilibrium of the NetApp during its operational state. That is to say that during the initial deployment the NetApp entails an initial state that can be disrupted. Initial state refers to the amount of resources that are allocated per component of the NetApp, the slice parameters that have been selected, the geographical constraints of the deployment, etc., the amount of instances per component. Each of these aspects serves a different functional purpose. However, based on activated policies, some of these aspects may be amended. Any change has to be reflected on respective control-plane signalling. The NetApp Lifecycle Manager is responsible for materializing these changes and **orchestrating all steps, in order for the new equilibrium to be maintained**.

Monitoring & Policy Execution

The Policy Execution module is responsible for the enforcement over the deployed NetApps following a continuous **match-resolve-act** approach. Specifically, the **match** phase regards the mapping of the set of applied rules which are satisfied based on the **data streams coming from the monitoring infrastructure**, the **resolve** phase regards the process of **conflict resolution -if any- among the satisfied rules,** taking into account the pre-defined salience of each rule, while the **act** phase regards the **provision of a set of suggested actions** to the orchestration module. Therefore, policies enforcement is realized through a **rule-based framework** that attempts to derive execution instructions based on the current set of data and the active rules; rules are associated with the deployed service graphs at each point in time.

The **rules engine** consists of (i) the **working memory** (WM); facts based on the provided data, (ii) the **production memory** (PM); set of defined rules, and (iii) an **inference engine** (IE) that supports **reasoning and conflict resolution** over the provided set of facts and rules, as well as triggering of the appropriate actions.

As already mentioned, data is fed to the Policy Execution through the Monitoring module that is responsible to collect data based on a set of active monitoring probes, as well as to support a set of data management operations (e.g., calculation of average values in specific time windows).

The definition of **rules per policy** will be provided by the Policy Editor based on the concepts represented in the Context model. A service graph may be associated with a set of policies; however, **only one can be active** during its deployment and execution time. Each **policy consists of a set of rules**. Each **rule consists** of the **a)**

expressions part, denoting a set of conditions to be met and **b) the actions part**, denoting actions upon the fulfilment of the conditions.

Expressions may regard custom metrics of a service graph or a component/microservice. An indicative expression is as follows: "if componentX.avg_cpu_usage is greater than 80%" and can be combined in various ways (and/or/grouping/subgrouping) with other expressions.

3.3 5G Telco Administrative Domain

OSS Northbound Proxy

Creating a proper slice and managing it all along the lifespan of a 5G-enabled application **cannot be performed without revealing the insights of the network dynamicity,** e.g., the position of the users, the allocated channels, the position of the micro-datacenters, etc. Such information is available only in the OSS/BSS layer of the telecommunication service provider. However, this information is highly valuable during the preparation and instantiation of the network slice instance (see phase a and b above).

There is **no de-facto API** that can expose in a structured format this type of information. However, several notable standardization efforts have been performed, in order to come up with a commonly acceptable API. 5G-INDUCE will not fall into the loophole of vendor lock-in as far as OSS/BSS API is concerned. Instead, it will perform the implementation using an abstracted open API allowing each vendor (of the use cases) to perform ad-hoc implementation of the services. In other words, 5G-INDUCE will maintain a solid Northbound API that will proxify internal slice management and slice configuration modules.

Metadata Synchronizer & Shared Repository

The role of the Shared Repository was analyzed above. During the formulation of the Slice Intent the service provider has to select (among others) a set of VNFs that will participate in the VNFFGs of the data plane. The knowledge of these VNFs should not be considered as granted since the NAO and the OSS are running under different administrative domains. The purpose of the shared repository is to act as a 'public' registry where each telco is announcing the VNFs/blueprints that can be supported during the slice creation.

OSS Core

The OSS Core is the main component that materializes the slices upon specific intents. It will be a radically enhanced version of the MATILDA OSS. In MATILDA, the interplay among NAO (termed VAO – Vertical Application Orchestrator in the MATILDA framework) and OSS has been realized through an intent-based interface, called 5G and edge computing slice negotiation interface. Upon NetApp deployment requests, it allowed the NAO to request, negotiate and obtain from the OSS both the needed computing resources at the edge facilities where to run NetApp components, and connectivity among such resources and User Equipment (UE). The OSS had the task of analyzing operational and performance (soft and hard) constraints expressed by the NAO slice request, and, consequently, of selecting the most suitable computing facilities and network services complying with the requirements.

The initial request produced by the NAO is called "slice intent" and contains the application graph annotated with QoS and operational requirements. Special links in the graphs represent the connectivity between front-end NetApp components and UEs (i.e., the PDU sessions to be realized by the radio mobile network and dedicated to the NetApp). When the NAO accepts the solution provided by the OSS, the OSS provides back the "materialized slice," which is the set of needed network services and resources instantiated and configured, ready to host the NetApp. MATILDA was able to fulfil these operations only at the instantiation and de-instantiation phases of a NetApp.

5G-INDUCE is extending the aforementioned 5G and edge computing slice negotiation interface for supporting the modification and reconfiguration of the slice within its lifetime and enabling advanced operations to deal with UE mobility and dynamic QoS/operation NetApp requirements.

In more detail, 5G-INDUCE NetApps (or some of their components) will be able to be lively scaled and relocated onto computing facilities in new geographical areas or at different network infrastructure aggregation levels in a transparent and smooth fashion. NetApp traffic from and to UEs will be steered accordingly (by coherently updating the configuration of network slices and related network services), while the platform will be able to adapt the NetApp geographical scope, and to properly scale the components and network services on each area depending on the number of hosted NetApps, the local workload, and the dynamic QoS/operational requirements.

For instance, live updates on the bandwidth/latency requirements of a link in the NetApp graph notified by the NAO might trigger the OSS to decide for scaling operation, reconfiguration, or and/or relocation on some of the involved VNFs. Updates on NetApp component scaling can trigger the OSS to interact with edge computing facilities to correctly re-dimension the resources allocated.

A second explanatory example is related to the geographical scope of a NetApp deployment, which is the geographical area covered by the radio mobile network where the NetApp will be made accessible from UEs with the requirements expressed in the slice intent. 5G-INDUCE will add the capability to the NAO of dynamically requesting a live change of the geographical scope of a running NetApp instance (i.e., for coping with UE mobility or for administratively extending/resizing the working area of the NetApp).

This geographical scope modification, on one side, entails the capability of the OSS to select/deselect proper resources at edge facilities updating/creating/deleting network services, and to dynamically reconfigure network slices to transparently/smoothly redirect UE incoming/outgoing traffic during the reconfiguration phases. On the other side, through a proper synchronization with the OSS, the NAO can update the NetApp instance graph by deploying/removing NetApp components in the selected edge facilities. In other words, the NAO will be equipped with a new function for the “right/left scaling” over the geographically distributed edge infrastructure.

A further main innovation point of the 5G-INDUCE OSS is that it will be extended to flexibly handle any NFV virtualization levels: network services can be composed by mixing VNFs realized not only with IaaS resources (as in MATILDA), but also with cloud-native containers over PaaS platforms (e.g., 3GPP 5G core network functions over Kubernetes clusters), and physical programmable/configurable devices (e.g., gNodeBs, eNodeBs, etc.). The number and the types of network services will be extended by including cloud-native 3GPP 5G network functions, and zero-touch procedures to enable the dynamic management of slices, traffic steering, etc.

In this respect, the 5G-INDUCE OSS includes a special module, namely, the NFV Convergence Layer (NFVCL), which fully drives NFV service orchestration along all the lifecycle phases. NFVCL communicates with an external NFV Orchestrator (e.g., ETSI Open-Source MANO – OSM) through standard ETSI NFV interfaces. During Day 0 operation, NFVCL produces and onboards onto the NFVO the ETSI SOL006 descriptors of services and of related Virtual/Physical/Container Network Functions by defining the needed number of virtual links and of virtual resources to be applied. In Day 1 operations, NFVCL requests the NFVO to instantiate network services selecting the computing facilities, and the networks where to attach network functions. At Day 2, NFVCL produces the configuration files and commands for each of the deployed VNFs and applies them through the VNF Managers at the NFVO.

Moreover, the OSS will be also provided with the capability to manage through a Metal-as-a-Service (MaaS) approach and to terraform bare-metal resources (i.e., servers, switches, routers, etc.) to install, to build, and to configure complex and distributed IaaS and PaaS environments, where to host VNFs and NetApps components.

The 5G-INDUCE platform will be specifically designed to adaptively exploit the programmability level offered by the underlying network infrastructure(s) by enabling the aforementioned capabilities when and where possible/necessary. More specifically, there are three different types of envisaged adaptations:

- **Full Programmability Adaptation:** it is the type of programmability where the control and data planes are created using solely 'internal API calls' on the terraformed infrastructure.
- **Partial Programmability Adaptation:** it is the type of programmability where the control and data planes are created using a combination of internal API calls and external ones that are treated as black boxes.
- **No Programmability Adaptation:** it is the type of adaptation where underlying resources are solely announced to the shared repository (cataloguing).

4 Requirements

This section presents the extracted requirements for the 5G-INDUCE platform and the modules under development. The section starts with the overall approach and the syntax used for the presentation of the functional requirements. The business-related (i.e., stakeholder-driven) requirements are presented first setting the basis for the functional requirements of the platform and modules that follow. The functional requirements are summarised in tables and split into generic and module-specific requirements. The key implementation remarks are extracted as a summary of the provided requirements. Finally, an analysis of the non-functional requirements is presented and linked to the functional ones.

4.1 Approach and Syntax

4.1.1 Overall approach

The 5G-INDUCE platform requirements are extracted with the goal to achieve a common integration plan among the various modules that meet the project objectives and the planned testing activities. Therefore, the main focus is on the **functional requirements** that are related to the implementation of the various modules at the application and network orchestration level, as well as with the interfacing of these modules towards a common integrated platform.

The adopted approach considered initially the extraction of the Business & Stakeholders' Requirements, in order to set the basis of the key functionalities that must be supported by the platform. This analysis started with the definition of the stakeholders and their role in the 5G-INDUCE ecosystem, and then moved to the requirements definition. The system-related functional requirements followed with the extraction of the basic platform requirements that meet the main project objectives and the initial architecture. Such requirements are generic and address the overall architecture providing the key definitions of the modules and their interconnection. As we progressed towards a more detailed definition of the platform's functionalities and the required modules, a more detailed set of requirements were defined also for the various modules including their targeted performance characteristics and technologies to be used for achieving a modular design. The final round of requirements is related to the revisiting of the overall architecture as one integrated system, thus focusing on the extraction of the interfacing among the various modules. This extracted set of requirements has defined in detail the reference architecture presented in the previous section and the initial implementation of the platform (delivered in D3.1 and described in D3.4).

The **non-functional requirements** have been also considered according to the ISO/IEC 25010 definition. The target, in this case, is to provide a more complete description of the platform's functional requirements with respect to a potential real deployment in the future. For this reason, the non-functional requirements are mapped to the functional ones, providing the required implementation features that should be considered for each case.

In parallel to the 5G-INDUCE platform requirements, the **use case requirements** have been studied and extracted presenting the supported functionalities, the end user needs, and the implementation requirements of the use cases. These requirements are included in D2.2 together with the detailed description of the use cases. The platform and system level requirements, presented in this deliverable, have taken into consideration the required use case functionalities and their deployment options.

4.1.2 Taxonomy of functional requirements

The functional requirements are split into the following three main categories:

- Business & Stakeholder Requirements (BSR): Includes the current / future requirements of the potential platform users (i.e., stakeholders) of the 5G-INDUCE solutions, with specific focus on the functions / features that are provided by the platform. These are further split according to the main stakeholders:
    - Industry (end user) requirements
    - NetApp developer requirements
    - Operator/Infrastructure owner requirements
- Generic Platform Requirements (GPR): Includes the overall system-related requirements and the interconnection requirements between different modules of the system. This is further split into:
    - GPR.A. Generic platform definition requirements – providing the functionalities of the main platform building blocks
    - GPR.B. Generic OSS implementation requirements – defining the key functionalities that must be supported by the OSS part of the platform
    - GPR.C. Generic NAO implementation requirements – defining the key functionalities that must be supported by the NAO part of the platform
    - GPR.D. User Interfacing requirements and information sharing – setting the key requirements for the interface of the platform with the end users
    - GPR.E. Generic security-related requirements – providing the security-related concerns for the platform modules (and not the security-related functions seen as NetApp components)
    - GPR.F. Generic services' use case related requirements – defining a generic set of NetApp functionalities that must be supported across the specific NetApp use cases
    - GPR.G. Data processing requirements and regulations – related to how the generated data streams and user information is handled by the platform, in compliance with directives, ethics rules and the DMP.
- Module Specific Requirements (MSR): Includes the specific requirements for the platform modules and the individual characteristics of the modules' components that implement the supported functionalities. This is further categorised according to the platform modules and their interfaces in a top-down approach from the end user interface to the network orchestrator:
    - NAO-to-end user interface
    - NAO modules
    - OSS-to-NAO interface
    - OSS modules
    - OSS-to-Network orchestrator interface
    - Network orchestrator modules

Each requirement is also characterised as a basic or extended requirement:

- Basic requirements are **mandatory** requirements for the proper function of the 5G-INDUCE platform and identified by the use of the word "*SHALL*" in their syntax.
- Extended requirements are **optional** requirements that add value to the basic platform functionalities and are compatible options planned for a future upgrade. These requirements are characterised by the use of the word "*SHOULD*" in their syntax.

4.1.3 Requirements Syntax

In order to have clear requirements we follow a specific type of syntax for all cases. The purpose is for each requirement to directly identify to what module or process the requirement refers to, what is the action required for this and what is the specific expected feature or value. There are three main types of requirements' syntax:

Type 1: [Condition] [Subject] - [Action] - [Object] [Constraint/Value]

- EXAMPLE: When signal x is received [Condition], the system [Subject] shall set [Action] the signal x received bit [Object] within 2 seconds [Constraint].

Type 2: [Condition] - [Action or Constraint] - [Value]

- EXAMPLE: At sea state 1 [Condition], the Radar System shall detect targets at ranges out to [Action or Constraint] 100 nautical miles [Value].

Type 3: [Subject] - [Action] - [Value]

- EXAMPLE: The Invoice System [Subject], shall display pending customer invoices [Action] in ascending order [Value] in which invoices are to be paid.

In the above types,

- The subject answers to 'Who/What' this requirement refers to.
- The object answers to 'what is the requirement for the subject'.
- The action shows 'in what way/how' this requirement is materialised.
- The condition provides the attributes that permit a requirement to be formulated, i.e. it makes a requirement for a specific subject dependent on pre-conditions rather than being generic for all potential cases.
- The constraints provide the restrictions for which the action will apply to the object.
- The Value answers to how much the entity of the requirement is (either measurable or comparable to something else).

In order to have one common and more generic syntax that is directly applicable to 5G-INDUCE modules, we decided to adopt the following structure that essentially combines the aforementioned types.

[**Subject**+(Condition)]	[**Action**]	[**Object**+{Constraint and/or Value}]

Other than the main Subject-Action-Object format, the specific syntax allows (where applicable) to state directly any pre-conditions for the module, function or interface, as well as the constraints to meet the requirements especially when targeted values are defined. This syntax accommodates the various types of requirements (i.e., definition, design, implementation, KPI, etc., type of requirement) and provides a modular approach that can be extended or defined in more detail by setting conditions and/or constraints to the targeted requirements.

4.2 Business related requirements

5G-INDUCE targets the development of an end-to-end orchestration platform, with the scope to enable advanced application-oriented NetApp deployment and management service solutions to the broader Industry 4.0 sector, which is recognised as one of the fastest growing and demanded vertical sectors in 5G [15]. The targeted business sector players bring their own market-oriented requirements on the introduced platform, according to their role, and with the scope to fulfil specific market potentials. In the following paragraphs, the 5G-INDUCE stakeholders are identified, while also their role in the overall 5G ecosystem, as well as their specific needs and market potentials, are presented. Furthermore, the business-driven requirements for the orchestration platform relevant to each stakeholder are listed.

4.2.1 Application software developers

Profile: The developers are primarily SMEs or Software development companies with expertise on end user applications. They develop software solutions that are typically custom-made to the specific needs of the customers (vertical industries) and with the goal to provide an end-to-end service solution or support specific customer-oriented functionalities. The targeted customers are either directly the vertical industries or other system integrators and service providers providing complete solutions to vertical industries.

Role in 5G-INDUCE: The developers provide the NetApp components for the hosted use cases (services). Their solutions may range from specific or general-purpose application components (to be adopted by third party NetApps), to complete NetApp solutions tailor-made for their customers' needs.

In 5G-INDUCE, application software developers provide the NetApps for the supported use cases, as well as additional NetApp components for adding security features.

Market potentials from 5G-INDUCE: 5G-INDUCE provides the enabling environment for developers to easily address the specific needs of end users, without having to deal with service providers' specific deployment platforms and processes. Essentially, this empowers SMEs to acquire a substantial market share in the industrial sector, which is currently monopolised by large service providers and software houses. In more detail, the key potentials include:

a) direct porting of services or targeted application components to an industry-oriented 5G ecosystem,
b) easy update of the offered services to keep track of the customer needs,
c) capability to exploit the underlay 5G programmability in the offered NetApps,
d) capability to greatly improve the quality of the offered solutions by building on top of other NetApps where applicable,
e) reusability of specific NetApp components for other cases thus minimising the overall development time.

Key requirements to bring about these potentials: The following table summarises the key functional requirements for the application software developers in order to maximise their market potentials.

Table 4-A: Business related requirements for application software developers

ID	Requirement		
	Subject + (Condition)	**Commitment**	**Action + Object + (constraint)**
BSR.D1	The NAO graph composer	SHOULD	provide automated component connection features based on the provided information for the declared components
BSR.D2	The NAO graph composer	SHOULD	update automatically the components' features based on the provided connectivity choices
BSR.D3	The NAO	SHOULD	provide alternative graphs for the same NetApp, allowing the selection of the required one based on the use case connectivity requirements
BSR.D4	A common metadata structure model	SHALL	be defined for easier and faster check of compatibility between available 3rd party components introduced
BSR.D5	A components' repository	SHOULD	be available to developers for selecting existing NetApp components for alternative NetApp deployments or new ones
BSR. D6	A components' repository	SHOULD	be made openly available through software (upon agreement of the related developers)
BSR.D7	NetApp validation and verification	SHALL	be made available through a DevOps testbed before launching this in an actual environment

4.2.2 Network operators and infrastructure owners

Profile: The operators of the telecom network infrastructure provide all the required transport network connectivity to the end-users. Within 5G and by following the latest market trends, the capabilities of this stakeholders' category are potentially extended to include core network processing capabilities. In principle, computation resources at the network core level can be offered by any data centre provider attached to the operator's infrastructure; however, for this profile in 5G-INDUCE it is assumed that the managed infrastructure resources include both the required network and computation resources.

Role in 5G-INDUCE: The operator of the network infrastructure is the one who will potentially accept and install the 5G-INDUCE platform solution over its infrastructure. The platform will be accepted as a functional extension of the network orchestrator and used in order to map application requests over the infrastructure resources in an automated manner.

Market potentials from 5G-INDUCE: The inclusion of the 5G-INDUCE platform solution by the operators and infrastructure owners enhances their market position, as it enables the move from the traditional infrastructure resource management to the service management and even service creation. In this context, the network slices created by the operator over the managed infrastructure can be tailored to the specific requests of their vertical industry customers, and even be adaptive to the overall network and computation resources. The key market-related potentials for operators and infrastructure owners include:

a) the capability to provide slices of their infrastructure tailored to the need of vertical industry customers,
b) the adaptation of the provided resources according to the needs of the running applications,
c) the offering of a service provisioning platform for NetApp development attached to the actual distributed resources and connectivity capabilities,
d) the potential provisioning of NetApp creation and management services as added value service to their industry customers (in addition to connectivity services),
e) the easy interoperability with other infrastructure owners and even service providers (enabling small MNOs and private infrastructure owners to coexist and offer their services to multiple end users).

Key requirements to bring about these potentials: The following table summarises the key functional requirements for the operators in order to maximise their market potentials.

Table 4-B: Business related requirements for Network operators and infrastructure owners

ID	Requirement		
	Subject + (Condition)	**Commitment**	**Action + Object + (constraint)**
BSR.O1	The operator's network orchestrator	SHALL	be allowed to handle its own NFV orchestration functions separately and autonomously from the 5G-INDUCE platform
BSR.O2	The operator's network orchestrator	SHALL	be allowed to handle NFVI resources separately, especially those related to its own infrastructure
BSR.O3	The operator's network orchestrator	SHALL	have view on the attached industrial/private network capabilities, when processing resources are required to be provided for the support of the NetApps
BSR.O4	The operator's network orchestrator	SHALL	expose its programmability capabilities to the 5G-INDUCE platform

REMARK: The network operators, due to their role as customers of the 5G-INDUCE platform, are strongly linked with the non-functional requirements. The specific functional requirements presented above are mostly linked to the compatibility characteristic, since this has a major effect in the design of the OSS and network orchestrator interface.

4.2.3 Vertical end users (Industries)

Profile: The vertical end users are the customers of the provided NetApp management infrastructure. The specific focus of 5G-INDUCE is on services targeting industry 4.0 type of end users although in principle the platform can be extended to any type of vertical end user. The vertical end users provide the NetApp specific functional requests for the deployment and management of their targeted NetApps. The end users can provide also their local (private part) of their infrastructure as part of the overall 5G NetApp deployment infrastructure over the public part of the network (provided by the operator).

Role in 5G-INDUCE: The main role of the vertical end users is the definition of specific use case NetApps. The industry end users collaborate closely with NetApp developers in order to design the targeted 5G-enebled NetApps. They also provide the experimentation facility infrastructure for the deployment and demonstration of the use cases, followed by the targeted performance indicators for the validation.

Market potentials from 5G-INDUCE: The 5G-INDUCE NetApp deployment platform enables the industry end users to design and deploy custom made 5G solutions with added value to their operations. This added value can be offered with respect to: i) increased productivity, by optimising specific processes across manufacturing sectors, ii) reduced downtime, by avoiding unwanted failures and process interruptions, or predicting on time potential failures and iii) increased product quality, by maintaining high working standards and identifying early any potential defects. An additional parameter is related to the personnel status and determined by level of safeness, security, and happiness within industrial environments. Some key 5G operations with high market potential are addressed and include:

a) increased amount of automation, aiming at boosting the productivity of daily operations and the efficiency of the production line,
b) highly optimized supply chains targeting reduced costs and higher availability,
c) greatly facilitated maintenance and repair processes for limited down time and increased productivity,
d) increased protection through smart surveillance of critical infrastructures.

Key requirements to bring about these potentials: The following table summarises the key functional requirements for the Industry end users, in order to maximise their market potentials.

Table 4-C: Business related requirements for Industry end users

ID	Requirement		
	Subject + (Condition)	**Commitment**	**Action + Object + (constraint)**
BSR.I1	The Industry end user for the deployment of a NetApp	SHALL	be unaware of the underlay network capabilities (or limitations) and provide only the targeted use case requirements
BSR.I2	The Industry end user for the deployment of a NetApp	SHALL	be notified if certain targeted requirements cannot be met

ID	Requirement		
	Subject + (Condition)	**Commitment**	**Action + Object + (constraint)**
BSR.I3	The Industry end user for the deployment of a NetApp	SHOULD	be offered alternative deployments in case that certain targeted requirements cannot be met
BSR.I4	NetApp components and graph repository	SHALL	be visible to IT personnel in industries and presented in a clear manner.
BSR.I5	The NAO through the user interface (supported by the OSS)	SHALL	enable IT personnel to reconfigure the NetApp requirements on demand
BSR.I6	The NAO through the user interface (supported by the OSS)	SHOULD	allow the direct extension of NetApp functionalities and the update of application components, without the need to reinitiate the service
BSR.I7	The NAO through the user interface (supported by the OSS)	SHOULD	be able to inform the end user about changes or failures in the network
BSR.I8	The NAO lifecycle management (supported by the OSS)	SHOULD	allow continuous service availability, through automated resiliency in case of failure, for ensuring at least 99,99% availability (or as defined by the end user)
BSR.I9	The NAO through the user interface dashboard	SHALL	offer a visual representation of metrics and information collected
BSR.I10	The Industry end user for the visualisation of metrics	SHALL	be able to select and deselect which metrics to be visible

REMARK: The table includes the generic industry requirements that are common to any use case and related to the design features and functionalities of the 5G-INDUCE platform. The specific use case requirements are presented in D2.2. Key non-functional requirements also apply for this case and relate mostly with the clarity of the provided end-user interface, and the performance related quality characteristics (efficiency, reliability and security).

4.2.4 System integrator and Service provider

The three types of stakeholders mentioned above can fully address the market value chain for the deployment of the 5G-INDUCE solution, with respect to the NetApp service development, management and deployment. The system integrator and service provider is a fourth stakeholder type that has the role of the intermediary between the end user and the infrastructure provider. This is a typical role in the networking market today due to the gap between the end-user service deployment and the infrastructure usage at the core and edge parts of the network – on the contrary, the cloud service provision model that runs on top of the telecom infrastructures offers a direct customer interface between the large cloud service providers and the registered end users.

The 5G-INDUCE solution can potentially eliminate this gap by making NetApps directly available to end users and by enabling the support for individual and custom-made application components through attached developers (mainly software solutions SMEs). As mentioned above, this model assumes an extended role for the operator, providing services to any attached end user. The important additional feature here is the integration of the end user infrastructures to the operator infrastructure for the provision of the required service, in full support of the requested end user requirements.

The realisation of the above vision requires first the full deployment of a user-friendly NetApp platform like 5G-INDUCE and the attraction of SME developers to provide NetApp solutions. The infrastructure owners are then ready to support such models, especially by enabling 5G stand-alone network services and utilising properly the slicing of their infrastructure to the end users. The last ingredient is the creation of a certain culture for the industry (and any vertical) end user that will trust its 5G operations to be executed in this 5G environment.

At this point, the role of the system integrator is important and comes to fill in the gap between the end user needs and the new NetApp capabilities offered by the infrastructure.

Profile: The system integrator provides a wholistic approach to the end user according to its needs. It also interfaces with the infrastructure owner (operator) to provide the connectivity and identifies (or provides on its own) the resources for the NetApp solutions (SMEs, developers, etc.).

Role in 5G-INDUCE: In 5G-INDUCE the system integrator is seen as an entity that has full knowledge of the 5G-INDUCE NAO user interface and its capabilities. This is used in order to translate the end user demands into NetApps by developing the required application graphs and communicating with the operator's infrastructure. The system integrator can also have the role of the 5G-INDUCE platform promotion to the operator.

Market potentials from 5G-INDUCE: Typically, the system integrators are targeting custom-made deployments from network infrastructure to software applications. A new potential is now created with 5G-INDUCE, in which the system integrator can obtain part of the big service provider role and offer NetApp development and also deployment and management services to a large number of end users, ranging from large enterprises to small ones requiring some automated processes with high throughput and latency guarantees. Another important potential is the support of collaborative services among different types of end users (e.g., manufacturing, shipping and logistics).

Key requirements to bring about these potentials: The system integrator has no specific requirements but combines the key requirements from all three stakeholders mentioned above.

4.3 Lists of functional (Technical) requirements

4.3.1 Generic system requirements

A list of the generic system functional requirements is presented below and is split into various categories that relate with the functional parts of the 5G-INDUCE platform

Table 4-D: Generic system requirements

ID	Requirement		
	Subject + (Condition)	**Commitment**	**Action + Object + (constraint)**
A. Generic platform definition requirements			
GPR.A1	The platform	SHALL	include the NAO and OSS as its basic module units
GPR.A2	The platform, through the OSS unit	SHALL	be aware of the available/detached Edge and Core Cloud domains for the deployment/maintenance of network functions and application components
GPR.A3	The platform, through the OSS unit	SHALL	be able to identify and reserve the necessary resources (computational and networking) for the deployment of NetApp (i.e., slice day-0 configuration)

ID	Requirement		
	Subject + (Condition)	**Commitment**	**Action + Object + (constraint)**
GPR.A3	The platform, through the OSS unit	SHOULD	support the reconfiguration of the resources (of already deployed NetApps) upon request from the NAO
GPR.A4	The platform, through the NAO unit	SHALL	provide the interface to the NetApp developers for the porting of their NetApp components independently of the underlying OSS layer
GPR.A5	The platform, through the NAO unit	SHALL	be able to manage the lifecycle of the onboarded applications
GPR.A6	The platform, through the NAO unit	SHALL	be able to provide SLAs, polices and high-level network requirements following the NetApp requests
GPR.A7	The platform, through the NAO unit	SHOULD	provide the capability to reconfigure the allocated resources (for already deployed NetApp) with the goal to maintain the requested SLAs
GPR.A8	5G-INDUCE platform	SHALL	be able to operate with 5G stand-alone infrastructures
GPR.A9	The 5G-INDUCE platform	SHOULD	have guaranteed availability, for critical components depending on the use case requirements
B. Generic OSS definition requirements			
GPR.B1	The OSS components	SHALL	be aware of the underlay programmable infrastructure including network, computational and radio access resources
GPR.B2	The OSS components	SHALL	be able to extract on demand the required set of resources from the application requests (coming from the NAO in the form of slice intents)
GPR.B3	The OSS components	SHALL	provide the allocation plans (main and alternative plans) for the NetApp deployment that best fit the NetApp requests
GPR.B4	The OSS components	SHALL	have compatibility with OSM NFVO/VNFM
GPR.B5	The OSS components	SHOULD	have compatibility with proprietary NFVO/VNFM (other than OSM), provided that these expose all their programmability aspects
GPR.B6	The OSS components	SHOULD	have compatibility with proprietary PNFMs and WAN RMs, provided that they expose documented programmable interfaces
C. Generic NAO definition requirements			
GPR.C1	The NAO components	SHALL	be defined and registered in a Docker Compose / Docker Format
GPR.C2	The NAO components	SHALL	claim the minimum execution requirements in terms of CPU, RAM, STORAGE
GPR.C3	The NAO components	SHALL	declare their exposed interfaces
GPR.C4	The NAO components	SHALL	declare their required interfaces with other components
GPR.C5	The NAO components	SHALL	declare the environmental variables the components use
GPR.C6	The NAO components	SHALL	be tagged with potential scaling features (scalable or not)

ID	Requirement		
	Subject + (Condition)	**Commitment**	**Action + Object + (constraint)**
GPR.C7	The NAO components	SHALL	claim for: GPU, Storage Persistence, Radio Connectivity, Deployment Location
GPR.C8	The NetApp links	SHALL	be tagged with the high-level network requirements (Minimum Throughput, Jitter, Packet Loss, Delay)
D. User Interfacing requirements and information sharing			
GPR.D1	The platform frontend interface (GUI)	SHALL	provide the necessary resource deployment information to the operator, via a monitoring dashboard
GPR.D2	The platform frontend interface (GUI)	SHALL	provide the NetApp deployment environment to the NetApp developers
GPR.D3	The platform frontend interface (GUI)	SHOULD	provide the necessary performance metrics to the NetApp owners (Industry)
GPR.D4	The platform	SHOULD	Support (E2E) performance monitoring of NetApps
GPR.D5	The platform	SHOULD	Provide (a certain level of) autonomy loop for DevOps (or any other NetApp stakeholders)
GPR.D6	The edge/MEC platform	SHOULD	Advertise available MEC services (RAN awareness, positioning services, etc.)
GPR.D7	The platform	SHOULD	Provide (or support) Security as a Service capabilities
E. Generic security related requirements			
GPR.E1	The platform	SHOULD	be able to use tools to automate configuration and deployment (minimise the human presence in the loop)
GPR.E2	The platform	SHOULD	be tolerable enough to survive failures and various attacks (resiliency).
GPR.E3	Connections among systems and equipment	SHALL	be secured by authentication and encryption
GPR.E4	The platform, through the NAO dashboard	SHOULD	be able to offer a more intuitive visualization of threats to cybersecurity professionals
F. Generic services' use case related requirements			
GPR.F1	Data transmission/reception (UL/DL)	SHOULD	be able to transmit/receive any kind of data
GPR.F2	Voice communication	SHOULD	support voice interactive communication in real time full duplex
GPR.F3	Video transmission/reception	SHOULD	be able to transmit/receive video information of different formats and definitions/resolutions and frame rates (i.e., 4K or 8K)
GPR.F4	Mobility	SHOULD	support the user need to receive and transmit information while moving
GPR.F5	Location Information	SHOULD	provide the location of the end-device or user (within a certain accuracy)

ID	Requirement		
	Subject + (Condition)	**Commitment**	**Action + Object + (constraint)**
GPR.F6	Low Latency	SHOULD	allow the required time between issuing a request or transmitting of a piece of information and receiving a response
GPR.F7	Edge Computing	SHOULD	provide computing power (CPU/GPU) at close proximity with the end device
GPR.F8	Edge Storage	SHOULD	store and retrieve a large amount of information with the least possible delay for storage and retrieval actions
GPR.F9	Reliability	SHOULD	indicate whether the Service is to be provided continuously
GPR.F10	Area Dependent Interactivity	SHOULD	support spatial distribution of the end-users/devices
GPR.F11	Privacy	SHOULD	guarantee protection level of usability and integrity of user data, equipment and network
GPR.F12	Traffic Type	SHOULD	specify traffic characteristics that the end-users (as a whole) receive and/or generate
G. Data processing requirements and regulations			
GPR.G1	Processing of data	SHALL	be aligned with the General Data Protection Regulation (EU) 2016/679 (GDPR)
GPR.G2	5G-INDUCE components and processes	SHOULD	follow the NIS directive
GPR.G3	5G-INDUCE components and processes	SHOULD	follow schemes established under the Cybersecurity Act
GPR.G4	The processing of personal data	SHOULD	be avoided in all cases
GPR.G5	The processed personal data, in case these are required for running a use case	SHALL	be restricted only to minimum necessary for fulfilling identified purposes
GPR.G6	The storage of personal data by the platform, in case these are required for running a use case	SHALL	be restricted to the minimum period necessary to fulfil the purpose of collection, as determined by each use case
GPR.G7	The platform modules	SHALL	be unable to store personal data outside the jurisdiction of the EU/EEA, (except where adequate data protection framework exists in a non-EU/EEA jurisdiction according to the GDPR).

4.3.2 Module specific requirements

The module specific requirements are listed in the following table providing the functional requirements related to the implementation of the platform modules. A top-down approach is followed with respect to the platform structure.

Table 4-E: Module specific requirements

ID	Requirement		
	Subject + (Condition)	**Commitment**	**Action + Object + (constraint)**
		A. NAO GUI	
MSR.A1	The NetApp Composer	SHALL	translate all the high-level requirements into metadata models capable for propagating all the appropriate information to the OSS
MSR.A2	The NetApp Composer	SHALL	provide a GUI for selecting and connecting the application components
MSR.A3	The NetApp Component Repo	SHALL	provide a catalogue of the available/onboarded application components
MSR.A4	The Policy Editor	SHALL	provide a GUI for setting the policies of either elasticity, security, or deployment
MSR.A5	The NetApp Graphs Repo	SHALL	provide a catalogue of the available/onboarded application graphs
MSR.A6	The NAO dashboard	SHALL	provide visual information on collected metrics
MSR.A7	The NAO dashboard	SHOULD	provide network status related information to infrastructure administrator
MSR.A8	The NAO GUI	SHALL	authenticate the registered end users according to their profile
MSR.A9	The NAU GUI	SHALL	restrict the actions of registered users according to their profile
MSR.A10	Logging and monitoring	SHOULD	be sufficient to protect against breach
MSR.A11	Authentication and sensor-derived data	SHALL	be safely encrypted
MSR.A12	User authentication levels and hierarchy	SHOULD	be properly and safely managed
		B. NAO	
MSR.B1	The NAO deployment manager	SHALL	provide the optimal only deployment plan
MSR.B2	The NAO execution Manager	SHALL	supervise the overall execution of the application
MSR.B3	The NAO monitoring Engine	SHALL	collect application-level metrics per component
MSR.B4	The NAO monitoring Engine	SHOULD	collect medium-related performance metrics
MSR.B5	The NAO Policy Engine	SHALL	provide inferences over acquired data and support runtime policies enforcement

ID	Requirement		
	Subject + (Condition)	**Commitment**	**Action + Object + (constraint)**
MSR.B6	The NAO Analytics / Profiling Engine	SHALL	produce advanced insights through machine learning mechanisms and provide real-time profiling of the deployed components, application graphs and VNF
MSR.B7	The NAO Service Discovery Mechanism	SHALL	provide registration and consumption of application-oriented services
C. NAO-OSS interface			
MSR.C1	The NAO Slice Manager Translator	SHALL	translate all the high-level requirements into metadata models capable of propagating all the appropriate information to the OSS
D. OSS			
MSR.D1	The OSS	SHALL	receive slice intent from the NAO and coordinate all the other control blocks in the 5G platform to set up 5G network slices and reserve edge and core cloud resources to host NetApp components
MSR.D2	The OSS	SHALL	process slice intent received from the NAO to materialize network services and reserve edge and core cloud resources that fit encoded NetApp requirements
MSR.D3	The OSS	SHALL	provide the NAO with a positive acknowledgement that includes descriptors of edge and core datacentres that have been selected to host NetApp components, (upon the reception of a slice intent that has been successfully processed)
MSR.D4	The OSS	SHOULD	interface the following 5G network control blocks: NFVO/VNFM, WAN RMs and VIMs
MSR.D5	The OSS	SHALL	support the deployment of NetApp functions as Docker containers over VMs
MSR.D6	The OSS	SHALL	support the deployment of NetApp functions as VMs
E. OSS-Network Orchestrator			
MSR.E1	The NFV Convergence Layer	SHALL	exploit standard ETSI NFV interfaces to communicate with OSM
MSR.E2	The NFV Convergence Layer	SHALL	produce and onboard the ETSI SOL006 service descriptors onto the NFVO
MSR.E3	The NFV Convergence Layer	SHALL	request the NFVO to instantiate network services selecting the computing facilities and the networks where to attach network functions
MSR.E4	The NFV Convergence Layer	SHALL	produce the configuration files and commands for each of the deployed VNFs, and apply them through the VNFMs at the NFVO
MSR.E5	The NFV Convergence Layer	SHOULD	communicate with other convergence layers interfacing NFVOs different from OSM.

ID	Requirement		
	Subject + (Condition)	**Commitment**	**Action + Object + (constraint)**
	F. Network Orchestrator		
MSR.F1	The NFVO/VNFM	SHOULD	manage the lifecycle of network services, by reserving, configuring and maintaining physical and virtual 5G network resources
MSR.F2	The WAN RM	SHOULD	manage and monitor the wide-area communication resources, to create network overlays to be used in a shared or isolated fashion by vertical applications and telecommunication services, and to provide information on which resources (e.g., VIMs, PNFs, etc.) can be selected in the distributed 5G infrastructure to create telecommunication services, in order to satisfy vertical application performance requirements
MSR.F3	The VIMs	SHOULD	abstract and expose computing, storage, and networking capabilities of edge and core datacenters within the 5G infrastructures
MSR.F4	VIMs	SHOULD	be materialized as OpenStack instances

4.4 Design and implementation remarks

The generic platform requirements (GPR.A) set the main focus of the 5G-INDUCE platform and split the development work into the NAO and OSS parts. Although the platform is a common integrated artifact, it is functionally split into the two modules. Individual requirements related to key design features of the OSS and NAO are also provided in GPR.B and GPR.C fields. The key functional requirements for the OSS are the awareness of the available/detached Edge and Core Cloud domains, the identification of the necessary resources (computational and networking) for the deployment of NetApp (required to identify the day-0 configurations), and the ability to reconfigure the resources (depending on the attached infrastructure properties and management capabilities). The key functional requirements for the NAO part are the proper interfacing with the NetApp developers for the porting of their NetApp components independently of the underlying OSS layer, the provision of the interfaces to manage the life-cycle of the onboarded applications and the provision of SLAs, polices and high-level network requirements following the NetApp requests combined with the capability to reconfigure the allocated resources (for already deployed NetApp) with the goal to maintain the requested SLAs. The above set the basis for the design and development of the individual functions in each module and determine the overall flow of information.

It is noted that the reconfigurability function is the most demanding, yet innovative, capability envisioned by the 5G-INDUCE platform and is declared as optional, in order to minimize the overall risk. Therefore, the effort will focus first on the provision of the key functions of the platform, required to demonstrate the NetApp onboarding and allocation capabilities and extended to the more complex reconfiguration functionality in the latest release of the platform.

Moreover, the generic requirements define the interfacing with the underlay network orchestration and infrastructure layers, denoting a direct compatibility with the OSM NFVO/VNFM as envisioned by the DevOps testbed. For the support of the experimentation facilities (and in general the exploitation of the solution towards any underlay orchestrator) it is defined that this will be supported through the appropriate interface, provided that the underlay infrastructure exposes all its programmability aspects and interfaces. This is an essential requirement for the integration of the 5G-INDUCE platform and the exploitation of its full potentials. Any limitation in the programmability will essentially limit the demonstrated capabilities, though without preventing the showcasing of the NetApp deployment and related use cases over fixed or predetermined allocations.

The generic user interfacing functional requirements are treated separately (GPR.D field) as they define key functionalities of the interface between the end users and the platform. The two main features are the NetApp deployment environment and the monitoring dashboard. The former is the platform interface through which the NetApp components are registered and onboarded and where the connectivity features are declared as end-user demands. The latter is the part that visualises the monitored features of the running applications and can be extended to any network-related resource representation. This part is closely related to the use cases' requirements defined in D2.2 and is open to innovative extensions with respect to the way that end users can communicate with the platform and the level of detail with which the initial application features and requests are introduced.

The security features are coming mainly as recommendations for the platform (GPR.E field) and involve three main topics. The first requires the authentication of both users and system components when they are connected. The second direction denotes the use of automation mechanisms that minimise the intervention of humans in the system. This relates mainly to the lifecycle management of NetApps and the automation capabilities introduced by the OSS layer. The third direction tackles the monitoring capabilities of the system and refers to the capability to inform the end users about potential threats. This can be further extended towards the creation of a common threat knowledge database, or even enable the automated deployment of enhanced security mechanisms when required and as a function to the NetApp performance characteristics set by the end users. It is noted that cybersecurity is not addressed directly in 5G-INDUCE. Instead, the project seeks the adoption of the generated knowhow from other activities, mainly in the form of added value VNFs provided as common NetApp components.

The services define a large set of practical requirements that identify key design features that must be supported by the 5G-INDUCE platform. In the GPR.F field there is a set of basic and generic features that enable the support of most of the use case services targeting the industrial manufacturing sector and the variety of use cases that can be deployed in such environments. These range from communication and networking features (e.g., supported data rate, latency, coverage reliability, etc.), to service type features (e.g., support of voice, video, mobility, traffic type separation) and to computational features (e.g., edge/cloud computing capabilities, storage). The requirements must be considered alongside the use case ones discussed in D2.2, which highlight also the detailed functional requirements for the supported use cases.

The last set of generic requirements (GPR.G field) relates to the use of collected and processed data and indicates the compliance with key European Commission directives and regulations dealing with data use, management and security. The main focus is on the handling of any potential personal data and the strict requirement for avoiding any related activity, including not only the storage and processing, but also the collection of such data.

The module-specific requirements (MSR) are extracted and presented per module following a top-down approach, from NAO end user interface to the network orchestrator. These requirements have defined the key design features of the 5G-INDUCE platform and the targeted interfaces. All the addressed features are discussed in detail in the presentation of the 5G-INDUCE architecture in section 3.

4.5 Non-Functional requirements

The non-functional requirements characterise the properties that the developed (i.e., functional) solutions must have, when the overall system is seen as an end product. More specifically, the non-functional requirements specify the criteria that can be used to judge the operation of a system, rather than specific behaviours of the implemented function, components, or modules. These requirements are essential for the product quality evaluation and therefore play an important role for the release of an end product at high TRL. According to ISO/IEC 25010, the quality characteristics are summarised in Figure 4.1.

Figure 4.1: Non-functional requirements (ISO/IEC 25010) [https://iso25000.com/images/figures/en/iso25010.png]

Here the non-functional requirements are introduced with the scope to determine which quality characteristics should be taken into account when evaluating the properties of the 5G-INDUCE platform solution, especially when this solution will move beyond the targeted TRL 7 at the end of the project.

In Table 4-F, we aim to provide a mapping between the functional requirements, and the non-functional ones, by highlighting the key quality characteristics that must be considered. The characterization and mapping are performed per group of functional requirements, since each group shares a number of functional requirements towards the same object. The highlighted (solid black) circle sign indicates the prime non-functional requirements that must be taken into consideration, while additional important requirements are marked too.

Each of the non-functional requirements are explained and discussed with respect to the 5G-INDUCE functional requirements in the paragraphs that follow Table 4-F.

Table 4-F: Mapping of functional to Non-functional requirements

List of 5G-INDUCE **Functional** requirements	Prevailing **Non-functional** Requirements (ISO/IEC 25010)							
	Functional Suitability	Performance Efficiency	Compatibility	Usability	Reliability	Security	Maintainability	Portability
GPR.A (Generic platform impl.)	○	○	○	○	○	○	○	○
GPR.B (Generic OSS implementation)	○	○	●		○			
GPR.C (Generic NAO implementation)	○	○		○	○			
GPR.D (User Interfacing and sharing)	○			●		●	○	○
GPR.E (Security capabilities)	○		○		●	●		
GPR.F (Generic services' use cases)	●					○	○	●
GPR.G (Data processing and regulations)	○		●			○		
MSR.A (End user NAO interface)	○			●		○	○	
MSR.B (NAO)	○	○		○	○			
MSR.C (NAO-OSS Interface)	○		○					○
MSR.D (OSS)	○	●	○		○			
MSR.E (OSS-Network Orch. interface)	○		●				○	○
MSR.F (Network Orchestrator)	○	○	○		○			○

Functional Suitability

According to [16], the functional suitability "represents the degree to which a product or system provides functions that meet stated and implied needs when used under specified conditions." In other words, it is the quality that rates the degree to which the implemented product meets the functional requirements set initially at the development phase and derived by the end user (i.e., targeted customer needs).

This characteristic provides a direct link to the functional requirements and therefore it is marked as a quality that applies for all functional requirement groups in 5G-INDUCE. It has particular meaning for the use cases' requirements highlighted here in the GPR.F field and analysed in more detail in deliverable D2.2.

Performance efficiency

According to [16], the performance efficiency "represents the performance relative to the amount of resources used under stated conditions." This quality characteristic relates mainly to the processing efficiency and the use of programming resources.

In 5G-INDUCE, the NAO and OSS are the main modules that require adequate performance efficiency, since they include intense processing functionalities. The related NAO part concerns the processing of the monitoring results in its analytics engine and the quick response to required reconfigurations over deployed NetApps. The processing requirements are more intense in the case of OSS and, in particular, in the efficient identification of the slice resources according to the NetApp requirements and the available solutions offered by the network. In addition, the timely response to reconfiguration requests is important. However, the OSS is deployed and executed in the core part of the network over the network orchestrator and with resources provided typically by the operator/infrastructure owner; thus, no strict requirements apply in terms of resource usage or deployment capacity.

Compatibility

According to [16], the compatibility refers to the "degree to which a product, system or component can exchange information with other products, systems or components, and/or perform its required functions while sharing the same hardware or software environment." For software solutions that apply in network and communication systems, this requirement translates into compatibility with common standards, including already released and established standards or planned standardisation activities for applied research-oriented solutions. A valid alternative is where the developed software operations are proven to be independent of any standards, e.g., through the use of compatible or adaptable programmable interfaces.

5G-INDUCE promises the development of an application-oriented NetApp deployment and management platform over any type of standardised network orchestrator. This sets a key functional requirement that denotes the development of the appropriate API towards the provided (by the operator) network orchestrator (see GPR.B4 and B5). A functional constraint to achieve this is that any interfaced proprietary network orchestrator exposes all its programmability aspects to the OSS. Furthermore, the project examines compatible solutions in the fields of slice specification, and service specification models, such as those defined by GSMA's Generic Slice Template model [17] and the TMF in [18] and [19], respectively. At the application level the compatibility requirements are more relaxed due to the diverse nature of the end user interfacing options. However, it is identified in 5G-INDUCE that by following a specified approach in the way that NetApp requests are passed to the platform can prove beneficial for the promotion of the platform over different types of vertical services; in this case, the TOSCA instance model [20] is considered. Finally, with respect to data management compliance standards, a key compatibility requirement arises to align any future high TRL solution to the specific EU regulations identified.

Usability

According to [16], the usability refers to the "degree to which a product or system can be used by specified users to achieve specified goals with effectiveness, efficiency and satisfaction in a specified context of use." The usability feature has particular meaning for solutions with end user interfaces or platforms that are destined to be adopted as services and integrated over existing infrastructures.

This characteristic is of paramount importance for the 5G-INDUCE platform's user interface and the capabilities it offers towards a) the NetApp developers for the proper boarding, update and reconfiguration of the application components, b) the vertical end users (potentially through the operator or a system integrator) for deploying and managing their NetApp, but also reconfiguring on demand the operational features, and c) the operator and end users for accessing the monitored information. Therefore, beyond the functional requirements of the NAO end user interface (provided in GPR.D and MSR.A fields) that fulfil all the targeted operation by the platform, the interface must additionally be easily understandable as well as operatable by different end users, provide the means to avoid critical human-oriented mistakes and provide a pleasing and satisfying environment to work with.

Reliability

According to [16], the reliability as a non-functional requirement refers to the "degree to which a system, product or component performs specified functions under specified conditions for a specified period of time." This quality characteristic relates with the ability of the system to be constantly operational and accessible and deliver its destined functionalities under normal operation. It also extends to the ability to quickly recover from faults and failures or equally to continue its operation when associated underlay hardware systems recover without the need for re-initiation and long reconfiguration processes.

In 5G-INDUCE the reliability feature is connected to the modules that perform the main platform operations (NAO and OSS). This can be achieved by contentiously keeping track of the module status and established configurations. The recovery, though, is in general a complex process that depends on the degree of system failures. The monitoring functions in the 5G-INDUCE platform have the additional goal to identify potential changes and resource failures and then act through the reconfiguration mechanism, thus providing a higher layer-oriented recovery. For a final high TRL, specific recovery mechanisms must be also considered for the NAO and OSS that allow quick recovery of data through updated images when this is required.

Security

According to [16], the security refers to the "degree to which a product or system protects information and data so that persons or other products or systems have the degree of data access appropriate to their types and levels of authorization." This is an important feature for any platform that hosts multiple end users and is prompt to potential attacks that target the interception of application data or the performance of malicious attacks to harm critical end user services.

In 5G-INDUCE, increased cybersecurity needs have been identified primarily for the end user interfacing modules and the handled NetApp datasets. Although the focus of the project is not on the establishment and enhancement of strong security protocols, user authentication and data encryption mechanisms are by default implemented in the NAO interfaces and the storage of the developed NetApps. Regarding the security at the NetApp dataset level, the overall service deployment concept relies on potential enhancements that can be deployed as NetApp components either on demand or in an automated manner, once a potential threat is identified (e.g., by a monitoring NetApp agent in the system). On the other hand, threats on the OSS module and network orchestration functions level are not likely to occur, since these modules are within the network operator environment and protected from unauthorised external access. Therefore, they are not directly addressed by the project. However, there are solutions from the project's consortium members that tackle security on the network functions level through risk assessment and enhanced attestation mechanisms and can be evaluated for potential high TRL releases beyond the project duration.

Maintainability

According to [16], the maintainability "represents the degree of effectiveness and efficiency with which a product or system can be modified to improve it, correct it or adapt it to changes in environment, and in requirements." The key feature in this quality characteristic is the ability of the software product or individual software components to adapt to any potential updates and changes, thus contributing to the longevity of the solution and consequently its market quality and presence.

The 5G-INDUCE platform solution follows a modular approach that enables parts of the overall platform to be easily updated, modified or even replaced without affecting other modules. The modularity applies also in the development of specific processes by following a micro-services deployment approach. Therefore, additional features, monitoring and analytic engines, optimisation processes and adopted data models can be easily tailored to targeted network needs and modified or updated accordingly.

There are two main fields identified within 5G-INDUCE where maintainability is required. The first relates to the underlay compatibility with the network orchestration platform. Here, the modular approach separates the OSS functionalities from the underlay orchestration layer through a specific API that communicates the programmability capabilities of the infrastructure for the optimum allocation of the slices. The second field relates to the adaptability of the 5G-INDUCE platform to different vertical application domains (beyond that of the targeted industry vertical). This requires adaptations only at the user interfacing part, in order to accommodate potential NetApp features that are not considered in the targeted project release, as well as end user tailored analytics and reconfiguration options; the reconfiguration, monitoring and visualisation engines remain though the same.

Portability

According to [16], the portability shows the "degree of effectiveness and efficiency with which a system, product or component can be transferred from one hardware, software or other operational or usage environment to another." The key feature of portability that mostly relates to networking environments is the adaptability of the offered solution to the underlay hardware (and supported software) infrastructure.

In 5G-INDUCE, this characteristic shares a common ground with the maintainability feature described above, since the adopted modular approach targets the adaptation of the platform to different infrastructure environments through easily modifiable interfaces. Indeed, the key targeted requirement here is to provide a platform solution that is seamlessly operable over diverse orchestration environments offered by different operators and infrastructure owners.

The portability is a key feature for the developed NetApps, which require to be easily adaptable and deployable at different end user sites. This feature affects the design of NetApps for the project use cases and is discussed in D2.2.

5 Conclusion

The current deliverable reports on the outcomes of the design phase of the project, focusing in particular on the implementation requirements for the 5G-INDUCE platform and the defined architecture.

More specifically the following key outcomes are reported:

- Definition of NetApp:
    - The NetApp is defined according to virtualizable functions (application and network) and by combining the application components with the requested networking functionalities.
    - The NetApps are linked to the functionalities of the platform modules.
- Extraction of reference architecture
    - The 5G-INDUCE architecture is defined in detail and spans across the NAO and OSS modules including the interfaces with the end users (in the northbound) and the network orchestrator (in the southbound)
    - An information flow (walkthrough) is described that highlights all the key functionalities of the platform. This is followed by detailed descriptions of the submodules
- Definition of requirements syntax and taxonomy
    - A clear syntax is defined for all type of requirements that denotes the object for which the requirement applies, while it allows an easy addition of constraints and targeted values
    - A requirements' taxonomy enables first the split into functional, non-functional and business-related requirements. Then, the functional requirements are further divided in groups according to the topic or part of the architecture that they apply.
- Business related requirements
    - The business-related requirements define the main needs of each stakeholder for the development of the targeted solutions. The key stakeholders identified are the application software developers, the network operators, the vertical end users, and the system integrator.
- Extraction of functional requirements
    - The generic platform requirements include the overall system related requirements and the interconnection requirements between different modules of the system. They extend to data usage requirements.
    - The module-specific requirements Include the specific requirements for the platform modules and the individual characteristics of the modules' components that implement the supported functionalities.
- Mapping of non-functional requirements to functional
    - The non-functional requirements follow the ISO/IEC 25010 model and determine which quality characteristics should be considered when evaluating the properties of the 5G-INDUCE platform solution. A mapping between the functional requirements per submodule and the non-functional requirements highlights which of these characteristics are of key importance for moving towards marketable solutions.

References

[1] 5GASP website: https://community.5gasp.eu/index.php/2021/07/12/what-is-a-netapp/.

[2] GSMA, "Network Slicing Use Cases Requirements", Jan. 2020, https://www.gsma.com/futurenetworks/wp-content/uploads/2020/01/2.1_Network-Slicing-Use-Case-Requirements-Booklet-1.pdf.

[3] GSMA, "From Vertical Industry Requirements to Network Slice Characteristics" Aug. 2018, https://www.gsma.com/futurenetworks/wp-content/uploads/2020/01/3.0_Characterisation-of-Network-Slices.pdf.

[4] TM Forum, "TMF IF1167 – ODA Functional Architecture", 2020.

[5] TM Forum, "TMF 633 – Service Catalog Management".

[6] TM Forum, "TMF 641 – Service Ordering Management".

[7] TM Forum, "TMF 638 - Service Inventory Management".

[8] TM Forum, "TMF909A – API Suite Specification for NaaS".

[9] MATILDA Project (2019) OSS for 5G-Ready Applications. https://private.matilda-5g.eu/Publications/DownloadPublication/22.

[10] 5GPPP ICT-41 projects: https://5g-ppp.eu/5g-ppp-phase-3-6-projects/.

[11] M. R. Raza, et al., "Reinforcement Learning for Slicing in a 5G Flexible RAN," in IEEE/OSA JLT, 2019.

[12] L. Zanzi, V. Sciancalepore, A. Garcia-Saavedra, H. D. Schotten and X. Costa-Pérez, "LACO: A Latency-Driven Network Slicing Orchestration in Beyond-5G Networks," IEEE Transactions on Wireless Communications, vol. 20, no. 1, pp. 667-682, Jan. 2021, doi: 10.1109/TWC.2020.3027963.

[13] Y. Zhang, et al., "Ml-based and qos-aware resource management for cloud microservices", in Proc, ASPLOS, 2021.

[14] ETSI-MEC, https://www.etsi.org/deliver/etsi_gr/MEC/001_099/027/02.01.01_60/gr_MEC027v020101p.pdf.

[15] 5GPPP White paper, "Empowering Vertical Industries through 5G Networks - Current Status and Future Trends", Aug. 2020.

[16] https://iso25000.com/index.php/en/iso-25000-standards/iso-25010.

[17] GSMA, Official Document NG.116, "Generic Network Slice Template", version 5.0, Jun. 2021.

[18] TM Forum, "TMF633 - Service Catalog Management API REST Specification".

[19] TM Forum, "TMF641 - Service Ordering Management API REST Specification".

[20] OASIS Open, "TOSCA-Instance-Model-v1.0-csd01", Nov. 2017.