



TNOVA

NETWORK FUNCTIONS AS-A-SERVICE
OVER VIRTUALISED INFRASTRUCTURES

GRANT AGREEMENT NO. 619520

Deliverable D2.1

System Use Cases and Requirements

Editor Jorge Carapinha (PTInS)

Contributors Aurora Ramos, Felicia Lobillo (ATOS), Thomas Pliakas (CLDST), Antonio Pietrabissa, Donato Macone, Francesco Delli Priscoli, Martina Panfili (CRAT), Yacine Rebahi (Fraunhofer), Jordi Ferrer Riera (i2CAT), Michael McGrath (INTEL), Paolo Comi (Italtel), Panagiotis Papadimitriou (LUH), Eleni Trouva, George Xilouris (NCSR), António Gamelas (PTInS), Dora Christofi, Georgios Dimosthenous (PTL), Georgios Gardikis (SPH), Athina Bourdena, Evangelos Markakis, Evangelos Pallis (TEIC), Mamadou Sidibe (VIO), Antonio Cimmino (ZHAW)

Version 1.0

Date June 15th, 2014

Distribution PUBLIC (PU)

Executive Summary

This report is the first public deliverable of the EU-FP7 Project Network T-NOVA, "Functions as-a-Service over Virtualised Infrastructures". It provides the results of Task 2.1 "System Use Cases and Requirements", which ran in the first five months of the Project as part of Work Package 2 (WP2) "System Specification".

The main goals of this deliverable are the definition of the basic T-NOVA use cases and the identification of the key system requirements. To accomplish these goals, a number of steps have been taken, including:

- Business analysis, including the definition of stakeholders, business roles and business scenarios. Four basic roles have been identified: Service Provider, Function Provider, Customer, Broker, as well as a number of additional optional roles.
- Definition of application scenarios intended to highlight the potential benefits of T-NOVA and to illustrate how the T-NOVA system could be used and exploited in practice.
- Analysis of the virtualized network functions to be integrated in the T-NOVA ecosystem and further developed in the scope of the project.
- Use case specification, in which 11 use cases describing the interactions between external actors and the system, based on business scenarios identified before, have been specified in detail.
- Requirements specification, based on the use cases defined in the previous step. The outcome has been a collection of 59 requirements addressing seven different areas: Management and Orchestration, Elasticity, Security, Resiliency, Service Continuity, Operations, Market / Commercial operability.

The present report provides the outcome of this effort. These results are expected to establish a common ground on which the remaining T-NOVA WP2 tasks (T2.2 to T2.6) and the other technical Work Packages (WP3 to WP6) will build their foundations.

Table of Contents

1. INTRODUCTION	5
1.1. MOTIVATION, OBJECTIVES AND SCOPE.....	5
1.2. DOCUMENT STRUCTURE	6
2. T-NOVA OVERVIEW	7
2.1. GOALS AND APPROACH	7
2.2. RELEVANCE TO ETSI NFV ISG	9
3. BUSINESS ROLES AND BUSINESS MODELS	11
3.1. INTRODUCTION.....	11
3.2. T-NOVA ROLES.....	11
3.3. BUSINESS MODELS AND BUSINESS STAKEHOLDERS.....	13
4. APPLICATION SCENARIOS.....	18
4.1. HIGH LEVEL SCENARIO	18
4.1.1. <i>Enterprise Version</i>	18
4.1.2. <i>Residential Version</i>	19
4.2. T-NOVA VNFs.....	19
4.2.1. <i>Virtual Security Appliance</i>	19
4.2.2. <i>Virtualised SBC</i>	21
4.2.3. <i>Virtualized DPI</i>	25
4.2.4. <i>Virtualized HG</i>	27
4.3. APPLICATION SCENARIOS WITH MORE THAN ONE VNF.....	30
4.3.1. <i>Enterprise Scenario: Attack against the SBC Component</i>	30
4.3.2. <i>Residential Scenario: vHG with DPI and security appliances</i>	31
4.4. RELATION TO ETSI NFV USE CASES	32
4.5. BENEFITS OF T-NOVA	33
5. USE CASES AND REQUIREMENTS	35
5.1. METHODOLOGY.....	35
5.2. USE CASES.....	36
5.2.1. <i>Basic Use Case diagram</i>	36
5.2.2. <i>Detailed use case description</i>	37
5.3. REQUIREMENTS.....	46
6. CONCLUSIONS	50
REFERENCES	51
LIST OF ACRONYMS	52
APPENDIX A. DETAILED REQUIREMENTS SPECIFICATION	54

Table of Figures

Figure 2-1 High-level visualisation of the T-NOVA architecture.....	8
Figure 2-2 Orchestration platform, services, and interfaces	9
Figure 3-1 T-NOVA roles.....	12
Figure 3-2 T-NOVA simplest business scenario.....	15
Figure 3-3 Broker stakeholder between Customer and Service Providers	16
Figure 3-4 Broker stakeholder between Service Provider and Function Providers.....	16
Figure 3-5 Broker stakeholder with several SPs and several FPs.....	17
Figure 4-1 Security appliance.....	20
Figure 4-2 High level model of an SBC	22
Figure 4-3 SBC interconnecting two sites	23
Figure 4-4 Usage of virtualized SBC video transcoding.....	24
Figure 4-5 DPI used for Monitoring and Statistics for Enterprise Customers.....	26
Figure 4-6 DPI used for traffic classification of multimedia streams	27
Figure 4-7 Home Gateway.....	28
Figure 4-8 HG fast enhancements.....	29
Figure 4-9 SA protecting SBC.....	31
Figure 4-10 vHG with DPI and security appliances.....	32
Figure 5-1 Requirement specification methodology outline.....	35
Figure 5-2 T-NOVA Overall Use case diagram	37

1. INTRODUCTION

1.1. Motivation, objectives and scope

T-NOVA, “Network Functions as-a-Service over Virtualised Infrastructures” is a European FP7 Large-scale Integrated Project, whose primary aim is the design and implementation of a management/orchestration framework for the automated provision, configuration, monitoring and optimisation of Network Functions-as-a-Service (NFaaS) over virtualised Network and IT infrastructures. T-NOVA leverages and enhances cloud management architectures for the elastic provision and (re-) allocation of IT resources hosting Network Functions. It also exploits and extends Software Defined Networking (SDN) platforms for efficient management of the network infrastructure.

The T-NOVA framework allows operators to deploy virtualised network functions, not only for their own needs, but also to offer them to their customers, as value-added services. Virtual network appliances (gateways, proxies, firewalls, transcoders, analysers etc.) can be provided on-demand “as-a-Service”, eliminating the need to acquire, install and maintain specialised hardware at customer premises.

Leveraging this NFaaS concept and in order to facilitate the involvement of diverse actors in the Network Function Virtualisation (NFV) scene as well as the attraction of new market entrants, T-NOVA introduces a novel concept of “NFV Marketplace”, in which network services and Functions offered by several developers can be published and brokered/traded. The NFV Marketplace enables customers to browse and select services and virtual appliances that best match their needs, as well as negotiate SLAs and be charged under various billing models. A novel business case for NFV is thus introduced and promoted.

T-NOVA activities are organised in 6 Work Packages, of which Work Package 2 (WP2) is focused on the definition of the overall system architecture, as well as per-subsystem specifications (Orchestrator Platform, Infrastructure Management, Marketplace, Functions).

This report is the first public deliverable of WP2 and summarises the results of the activities carried out in the scope of Task 2.1 “System Use Cases and Requirements”, which intends to establish a common ground across all T-NOVA work packages, by providing:

- A number of representative application scenarios built around the T-NOVA ecosystem.
- The definition of the Use Cases foreseen for T-NOVA system.
- The identification of T-NOVA system business roles and value chain, considering their potential interest in the adoption of T-NOVA.
- A collection of system requirements considering all the business actors and identification of their role in the T-NOVA ecosystem.

The approach adopted in the definition of roles, use cases and requirements followed best practices in requirements engineering. The definition of general scenarios was

the first activity carried out by Task 2.1. Based on the four network functions to be developed in the scope of the project (i.e. Virtual Security Appliance, Virtualised SBC, Virtualised DPI and Virtualised HG), these scenarios have been put forward by different partners to provide a representative spectrum of experience regarding the covered areas and provide the starting point for the subsequent phases of the work.

The results included in this report reflect the work carried out in the initial months of T-NOVA. As a result of the iterative strategy followed by the project, it is expected that the feedback generated by Work Packages 3-6, as well as remaining WP2 Tasks (2.2 to 2.6), will require revisiting some of the requirements specified in this document, or defining new ones, at a later stage in the project. If such changes are considered substantial, they will be included in a future second release of this document.

During the process of specifying requirements, a preliminary identification of system components, needed to carry out specific parts of the system, was performed. However, it should be noted that a description of the overall architecture and the definition of architectural components are only to be found in forthcoming Deliverable D2.21, therefore out of the scope of the present deliverable.

1.2. Document structure

Following this introductory section, the remaining part of the document is structured as follows:

- Section 2 provides a general overview of the objectives T-NOVA and clarifies the positioning of the project in relation to ETSI NFV, currently the most important standardization effort in this area.
- Section 3 addresses the basic T-NOVA business models and roles, including a description of their main business relationships.
- Section 4 presents a number of application / business scenarios that illustrate how the T-NOVA system would be deployed and used in a real environment. These scenarios show how the four virtualised network functions to be developed during the course of the project would fit into the T-NOVA landscape.
- Section 5 describes use cases supposed to cover the interaction between roles and the T-NOVA ecosystem in different stages of the service lifecycle, and collects the corresponding requirements. For the sake of readability, the full set of 59 requirements has been moved to Appendix A and only general conclusions are provided here.
- Finally, section 6 provides general conclusions and guidelines on how the results included in this deliverable will be further advanced in forthcoming stages of the project activities, not only within the scope of WP2, but other WPs as well.

2. T-NOVA OVERVIEW

This section provides a general overview of the objectives of the T-NOVA project and a preliminary description of the technical approach to be further developed in subsequent stages of the project. Further information on T-NOVA architecture and architectural components will be provided in future T-NOVA deliverables, namely D2.21, D2.31 and D2.41.

2.1. Goals and Approach

Network Functions Virtualisation is a concept aimed at virtualising network functions such as gateways, proxies, firewalls and transcoders, traditionally carried out by specialised hardware devices, and migrating those functions to software-based appliances, deployed on top of commodity IT infrastructure.

The migration of most of the in-network operations from hardware to software modules leads to various benefits including:

- Efficient management of hardware resources,
- Rapid introduction of new network functions to the market,
- Easy upgrade and maintenance,
- Exploitation of existing virtualisation and cloud management technologies for the NFVs,
- Significant CAPEX and OPEX reduction.

While software based versions of network functions have been available for a number of years, automation and deployment of these functions in a carrier grade environment have not received enough attention from the research and industry communities. To address these issues, the EU FP7 T-NOVA project suggests an integrated architecture allowing network operators not only to deploy virtualised Network Functions (NFs) for their own needs, but also to offer them to their customers, as value-added services (Network Functions as-a-Service, NFaaS). With T-NOVA, virtual network appliances (gateways, proxies, firewalls, transcoders, analysers etc.) can be provided on-demand "as-a-Service", eliminating the need to acquire, install and maintain specialised hardware at customers' premises. This dynamic provisioning of virtual network functions is achieved by means of an innovative "NFV Marketplace" where network services and functions created by a variety of developers can be published, acquired and instantiated on-demand.

Definition of T-NOVA architecture will be carried out in Task 2.2 and described in Deliverable D2.21 (and later refined in D2.22), therefore a detailed description of the architecture and architectural components is out of the scope of this document. A preliminary high-level view of the main T-NOVA building blocks is depicted in Figure 2-1. The system to be designed and developed has the objective of implementing all the functionalities (virtualisation, orchestration, and resource management) of a complete Network Function Virtualisation Infrastructure (NFVI), as defined by ETSI NFV (1).

In addition to the core NFVI functionalities, the T-NOVA system also provides an innovative “Network Function Store”, imitating the paradigm of the successful OS-specific “App Stores” for smartphones and tablets. To be more specific, the Network Function Store will facilitate the involvement of diverse actors in the NFV scene. It will enable for example third party developers to implement NFs and publish them as independent entities, accompanied with the necessary metadata. The NF Store allows customers to select the virtual appliances that are more appropriate to their requirements, “plug” them into their existing connectivity services and configure/adapt them according to their needs. Service request and initiation is carried out via a customisable front-end/brokerage platform. With the NF Store, T-NOVA introduces the NFV Marketplace, which enables new business cases and market opportunities, both to existing players and new entrants.

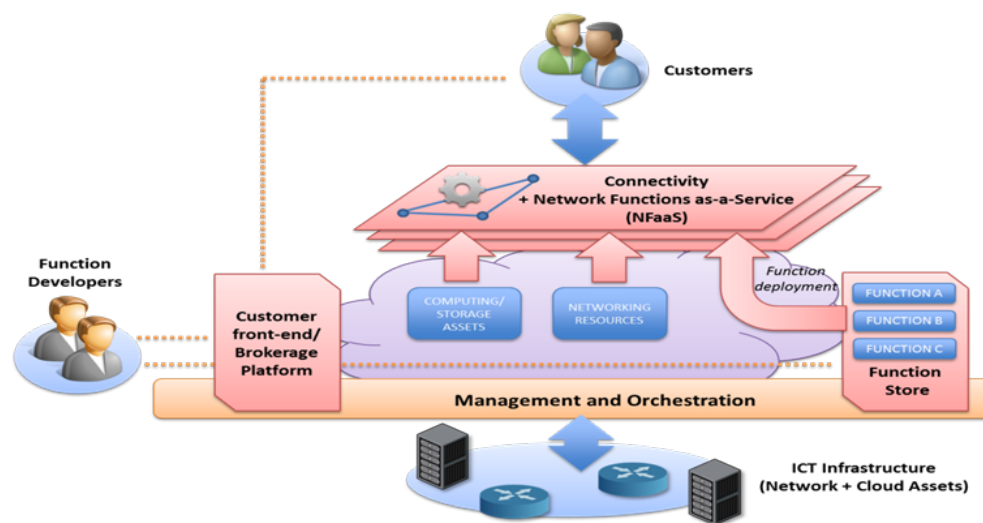


Figure 2-1 High-level visualisation of the T-NOVA architecture

In this context, the Management and Orchestration layer in the T-NOVA architecture plays a crucial role, as it addresses two critical issues in Network Function Virtualisation: (a) automated deployment and configuration of NFs, and (b) federated management and optimisation of networking and IT resources for NF accommodation.

As shown in Figure 2-2, the Management and Orchestration layer (shortly Orchestrator) acts as a middleware able to deploy and monitor T-NOVA Services by jointly managing Wide-Area Network (WAN) network resources and in-network cloud (compute/storage) assets. The Orchestrator is the highest-level infrastructure management entity, which orchestrates network and IT assets in order to compose and provision the T-NOVA services. It manages virtual network setup, traffic steering, Virtual Network Functions (VNF) instantiation and placement, and supervises/controls the provisioned service.

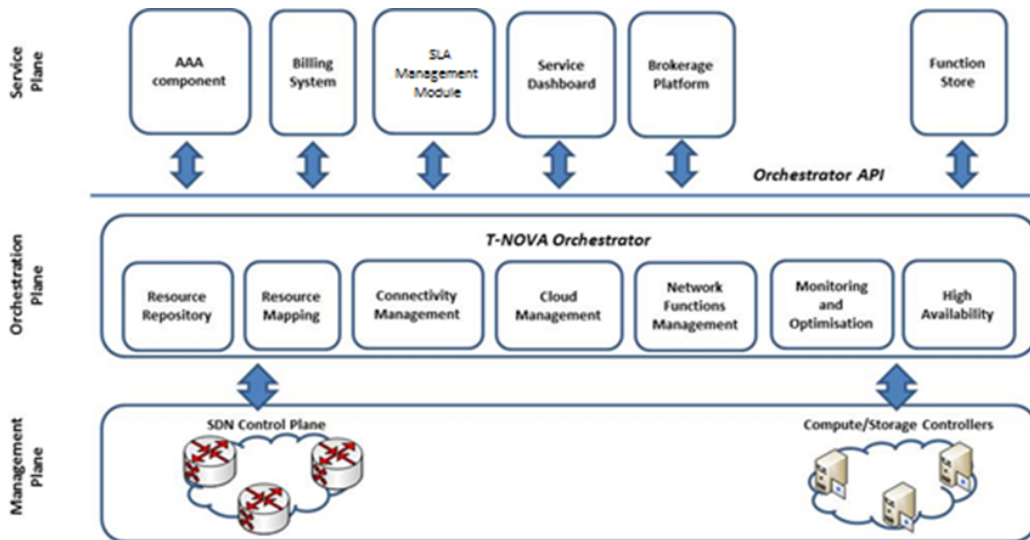


Figure 2-2 Orchestration platform, services, and interfaces

Details on the specific functionalities to be executed by each of these building blocks will be provided in future T-NOVA deliverables identified above.

2.2. Relevance to ETSI NFV ISG

In the area of Network Functions Virtualisation, ETSI NFV ISG undoubtedly represents the most relevant standardisation initiative. ETSI NFV ISG has published a set of high-level specifications in order to present the NFV ISG view and also to provide guidance to the telco industry. The documents that are publicly available through the NFV ISG information portal (2) are:

- NFV Use Cases document: Describes initial fields of application (3).
- NFV Requirements document: Describes the high level business and technical requirements for an NFV framework including service models (4).
- NFV Architectural Framework document: Describes the high-level functional architecture and design philosophy for an NFV enabling architecture, the virtualised network functions and the underlying virtualisation infrastructure (1).
- NFV Terminology document: A common repository for terms used within the NFV ISG documents (5).
- NFV Proof of Concept Framework document. The NFV ISG has launched a global call for multi-party NFV Proof of Concepts (PoC) to validate NFV approaches and to encourage progress towards interoperability and development of an open ecosystem (6). The document defines a framework for use within ETSI NFV to coordinate and promote this process.

Within the context of T-NOVA, these ETSI documents are considered to be key references, as the intention of the project is to align its approach to the current view of the NFV ISG. In this deliverable, the proposed NFV ISG terminology has been adopted in order to achieve consistency. Moreover, new terms have been introduced in order to cover aspects that the NFV ISG does not currently address. The high level

requirements set out by the ISG have been analysed and are taken as input on section 5. More technical functional and non-functional requirements related to the main components of the T-NOVA platform will be provided in the forthcoming architecture specific deliverables.

With regard to Use Cases (UC), the high-level UCs for the T-NOVA system are described within this document. These UCs describe the main interactions between the system and the actors. The ISG document adopts a slightly different approach, as it identifies Use Cases based on high-level usage and deployment scenarios. The relationship between the T-NOVA approach and the ISG Use Cases is discussed later in this document.

As a conclusion, T-NOVA scope is very closely related to the proposed view of the ETSI NFV ISG and the participating industrial partners will exploit the possibility of contributing to the second phase of the ISG lifecycle.

3. BUSINESS ROLES AND BUSINESS MODELS

3.1. Introduction

This section defines the T-NOVA roles and introduces the various business models that may arise from them, depending on the manner these roles will be implemented by different stakeholders in various scenarios. A stakeholder or actor can be defined as an individual, group of people, organisation or other entity that has a direct or indirect interest (or stake) in a system. Meanwhile, the correspondence between the roles played by specific stakeholders and business entities (e.g. users, network operators, service providers) is not necessarily 1-to-1. In fact, the same business entity can play more than one role. Role analysis is specifically focused on functionality (described in section 3.2) and not on who in practice plays those functions, which is introduced in section 3.3. The various potential business scenarios for T-NOVA will be examined in detail in WP8. Based on the same set of roles, multiple business models will be built in D8.12.

3.2. T-NOVA Roles

Several roles are involved in the T-NOVA value chain. Some of these roles may be grouped to be implemented by the same stakeholder depending on the different scenarios. However, we prefer to be exhaustive in the role analysis process in order to explore at a later stage as many different commercial relationships as the T-NOVA paradigm has potential to support. The roles identified in T-NOVA system (see Figure 3-1) are as follows:

- **End User (EU):** This is the end consumer of the purchased service, which is acquired by the Customer (C).
- **Customer (C):** The T-NOVA Customer who purchases T-NOVA services.
- **Service Provider (SP):** The SP provides a finished product to end customers. Services offered to end customers can be single network functions, or bundles containing a combination of functions from different Function Providers (FP), or a complete end-to-end network service (5).
- **Function Provider (FP):** The FP supplies virtual network appliances (gateways, proxies, firewalls, transcoders, etc.) eliminating the need for the customer to acquire install and maintain specialised hardware.
- **Broker (B):** The broker role performs trading between the customer and the service providers and between service providers and function providers. The broker fetches offerings matching the customer requirements and, depending on the applicable trading-policies, carries out the necessary actions for the

customer, the SP and the FP to agree on definite SLA¹s and prices to be applied.

- **Cloud Infrastructure Provider (CIP):** The CIP provides the cloud infrastructure where the NF will run on.
- **Network Infrastructure Provider (NIP):** The NIP provides the physical connection to the cloud infrastructure.
- **Service Integrator (SI):** The SI matches the suppliers providing the substrate for running the virtualised functions for the SP. Depending on the function features and requirements (SLAs included) and taking into account the different available possibilities, the SI makes the match so that the service can finally be delivered.

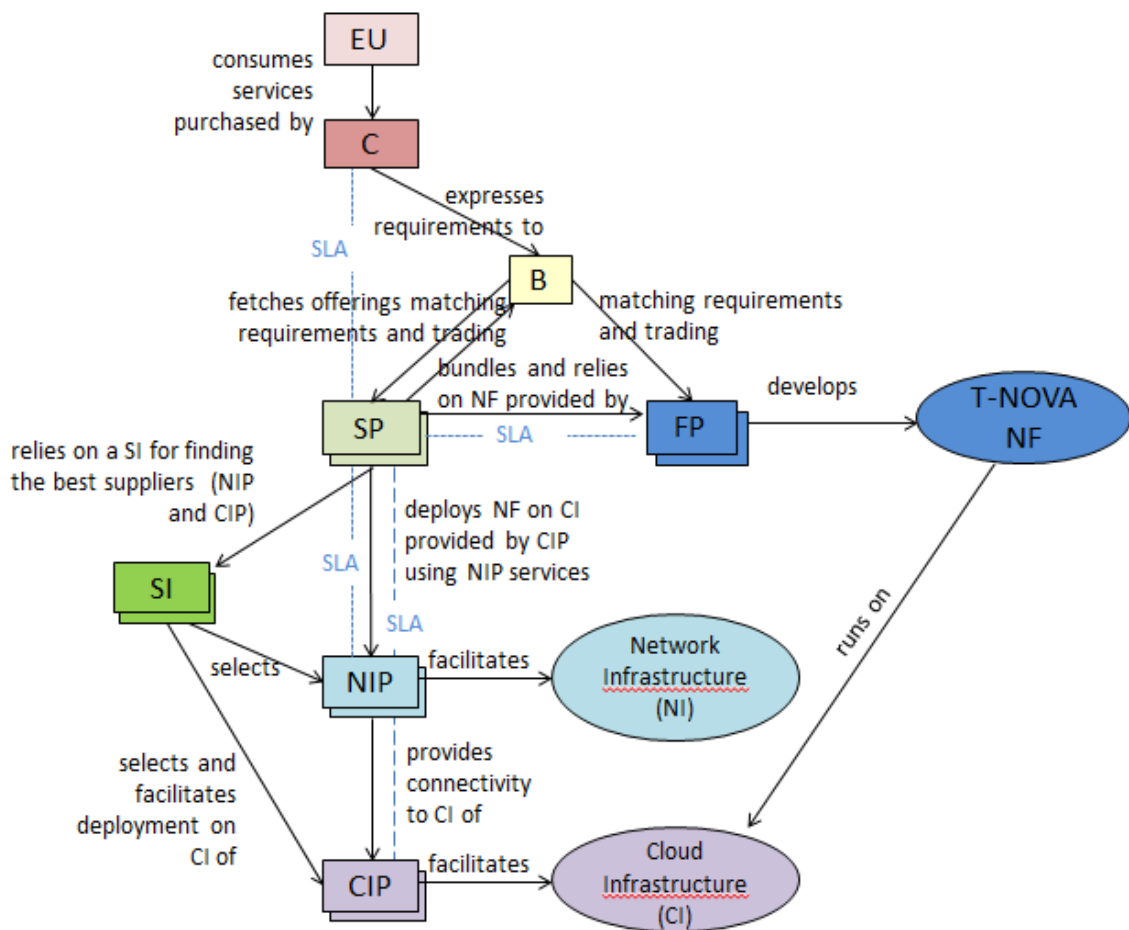


Figure 3-1 T-NOVA roles

¹ Service Level Agreement (SLA) is a negotiated agreement between two or more parties, recording a common understanding about the service and/or service behaviour (e.g. availability, performance, service continuity, responsiveness to anomalies, security, service ability, operation) offered by one party to another, and the measurable target values characterizing the level of services. SLAs always follow the same steps: publication, discovery, negotiation, provision, monitoring and evaluation. The SLAs can be customer SLAs (between the customer and the SP) and/or supplier SLAs between the SP and third suppliers (FP/IP). The evaluation of an SLA might produce a billable item (e.g. compensation fee, reward discount). NOTE: The scope of the above definition does not include business aspects of the SLA (9).

3.3. Business Models and Business Stakeholders

Different business scenarios can arise from T-NOVA value chain depending on how the T-NOVA roles are played by different business entities (stakeholders). In this subsection, the business interests and relationships for the roles previously defined are highlighted. We also explore the different commercial relationships that the T-NOVA paradigm will bring, since in practice not all of the T-NOVA roles will be implemented by different stakeholders, depending on the circumstances.

In the scenarios where each role is played by a different stakeholder, that is, there is a 1-to-1 relation, as represented in Figure 3-1, the general business interests and relationship of T-NOVA roles will be as follows:

- **End User (EU):** EUs will be in business scenarios such as a SP being a reseller. In this case, the EU would pay for services provided by the SP. In other contexts, this actor will not be present.
- **Customer (C):** The customer purchases and pays for a service provided by the SP. The customer might be involved in bargaining or auctioning processes in order to negotiate a final provider and terms (price and SLA) for the service purchased. The Customer agrees on a SLA with the SP. Whether or not the SLA has been fulfilled might produce a billable item that will later appear on the bill. This means that in the case where a service has not been delivered with the expected and agreed quality, the customer may receive a discount on their bill. The penalties for the faulting parties have to be agreed clearly within the SLAs.
- **Broker (B):** The broker business model as third party might be to be granted a commission per finalised trade and, thus, shares revenue with SPs or FPs. A broker might prioritise offerings depending on the business terms negotiated with the SP or FP.
- **Service Provider (SP):** The SP provides the service to the customer making alliances with FPs and with the B, sharing revenue with both. The SP might incentive the B in order to push certain offerings over others, even over competitors' ones. FP might receive their share from the SP in different billing modalities, which most likely will depend on how the offering is priced. For example, when a monthly fee is applied to an offering, the FP might also be paid monthly. For offerings involving more occasional services, the payments will also be prompt. The SP agrees on an SLA with FP. The result of this SLA, i.e. whether the SLA has been met or not, might produce a billable item that will later appear on the bill. This means that in case a service has not been delivered with the full expected and agreed quality, the SP might receive a discount on his bill. The penalties for the faulting parties have to be agreed clearly on the SLAs. These SLAs between the SP and FP have a relationship with the SLAs established between the C and the SP.
- **Function Provider (FP):** The FP is interested in providing as many functions as possible or in having their functions used or purchased as much as possible. The FP commercialises NFs through the Broker to the T-NOVA SP.

The relationship with the SP and B is a revenue sharing scenario, since the SP has to grant a share to the FP for the NF included in a service.

- **Service Integrator (SI):** The SI's customer is the SP. Among different SI, the SP might choose the most trusted one, or the SP offering the best conditions (price, etc.) for the same service. This stakeholder is not really fundamental in T-NOVA and we do not foresee that the T-NOVA paradigm supports great dynamism in these commercial transactions. Normally a unique SI would exist for a SP and this role would be played by the SP themselves.
- **Cloud Infrastructure Provider (CIP):** A CIP has a commercial relationship with the SI (in case it exists) or, otherwise, with the SP (in the case where they are not played by the same business entity, e.g. a Cloud Service Provider (CSP)).
- **Network Infrastructure Provider (NIP):** A NIP has a commercial relationship with the SI (in case it exists) or, otherwise, with the SP (in case they are not played by the same business entity, e.g. a CSP).

From this study of business cases it can be concluded that in context of T-NOVA there are some T-NOVA roles that will necessarily be played by different stakeholders, which we have called basic stakeholders, in the sense that they are likely to be part of all the use cases. The remaining stakeholders may or may not be included as appropriate.

Table 3-1 summarises basic and non-basic stakeholders:

Table 3-1 T-NOVA Stakeholders

Stakeholder name	Comment
Service Provider (SP)	Basic
Function Provider (FP)	Basic
Customer	Basic
Broker	Basic if the business scenario has several SPs. Optional if there is only one SP. The SP can contract or not a third party to perform trading among FPs to purchase VNFs. If the broker is not contracted, the SP will perform itself the trading among different FPs.
Service Integrator	Optional (likely to be a function played by the T-NOVA SP)
Network Infrastructure Provider	Optional (likely to be a function played by the T-NOVA SP)
Cloud Infrastructure Provider	Optional (likely to be a function played by the T-NOVA SP)
End User	Optional (may be needed only in some specific scenarios)

The most simplified version of the business relation landscape will be that in which there is only one SP and one FP, so the T-NOVA SP can play the broker role itself to trade its services, and the FP can trade their own NFs (unless the SP and FP prefer to contract a broker stakeholder to do it). The T-NOVA SP may also act in the role of the CIP and the NIP. This would be the case of a network operator who provides NFs over their own infrastructure (both cloud and network resources). Also, the role of the SI is most likely played by the SP. In residential scenarios, the C and the EU will be played by the same business entity that will be the T-NOVA Customer. The simplest basic scenario is represented in Figure 3-2.

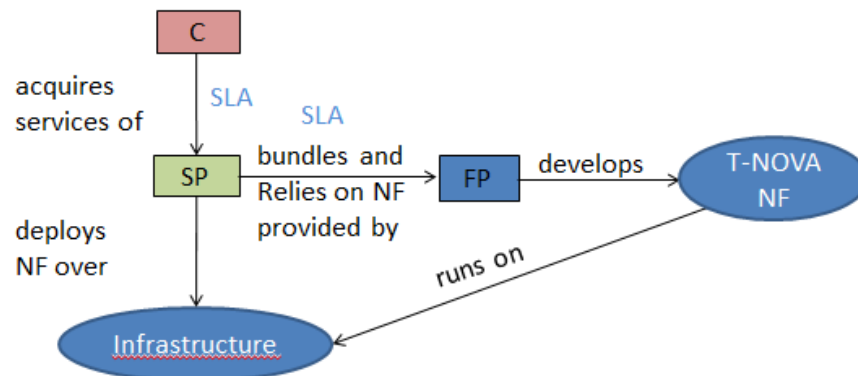


Figure 3-2 T-NOVA simplest business scenario

In this scenario, several SPs have accessed to the T-NOVA system to provide services. The broker role must be played by a separate stakeholder or third party entity, in order to offer the customer the best price options in the T-NOVA marketplace as a result of the trading among different SPs. In this case, the broker stakeholder being an intermediary player selects the offerings and conditions that match the customer requirements considering all the services provided by all the SPs in T-NOVA. In this business situation, as said before, the broker role is not part of the T-NOVA SP, but a third party that promotes dynamism and concurrence among the different SPs implementing trading mechanisms. The broker business model based around a third party scenario may be operated on a commission per finalised trade and, thus, shares revenue with SPs. A broker might prioritise offerings, depending on the business terms negotiated with a SP. This situation is represented in Figure 3-3.

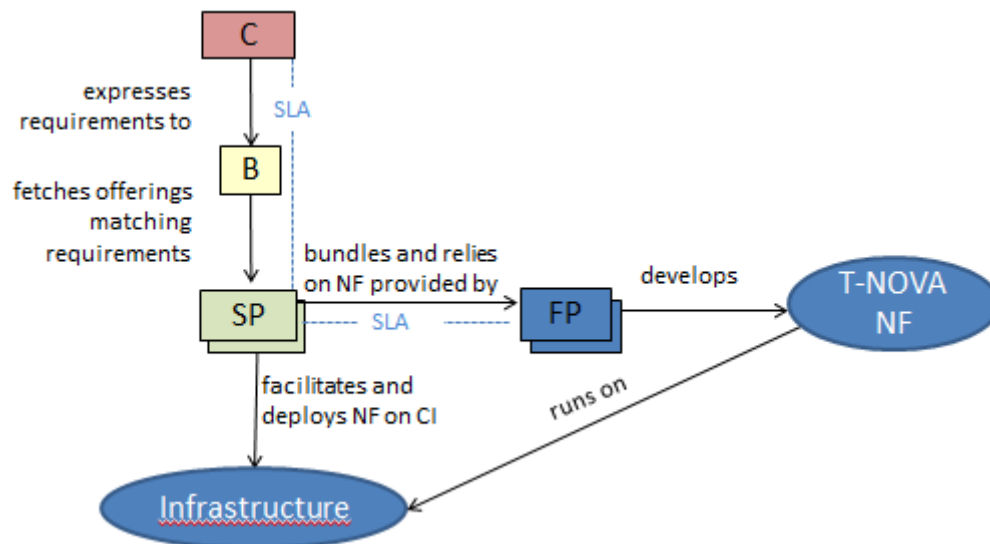


Figure 3-3 Broker stakeholder between Customer and Service Providers

Under normal circumstances, several FPs have accessed T-NOVA system to provide their VNFs. A SP will have the ability to decide if they want to contract a broker stakeholder to play the broker role among FPs, or if the SP itself looks for NFs to purchase in order to compose their own service offerings based on the combination of the best price according to the SLA and the customer requirements. In this situation, the broker is an intermediary player that provides a service to the SP performing trading among multiple FPs. Under this common scenario, the broker is not part of the T-NOVA SP, but rather a third party that promotes dynamism and concurrence among the different FPs implementing the trading mechanisms. The broker's business model as a third party might be based on receiving a commission per finalised trade, which is essentially a shared revenue model with the FPs. A broker might prioritise NFs depending on the business terms negotiated with each FP. This situation is represented in Figure 3-4.

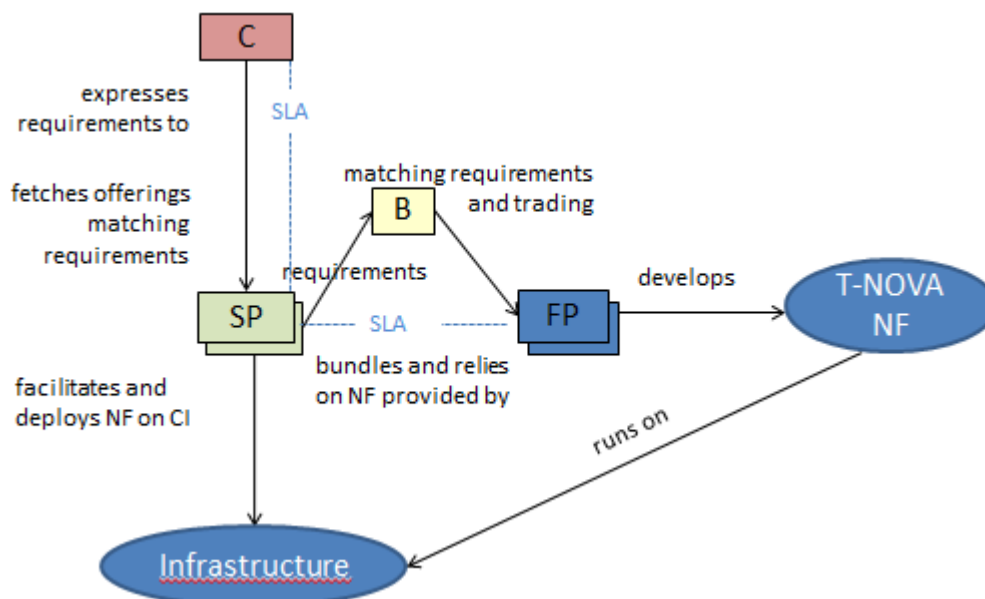


Figure 3-4 Broker stakeholder between Service Provider and Function Providers

In T-NOVA the most interesting and common scenario from the business point of view will be the one represented in Figure 3-5, where a broker stakeholder would perform the trading among several SPs and several FPs offering the customer the best price option considering their requirements.

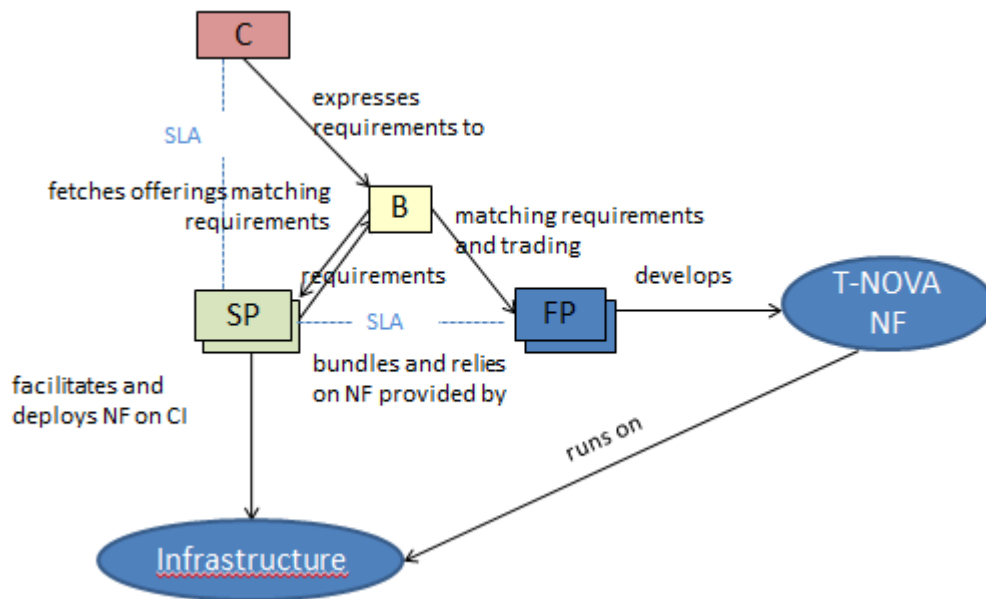


Figure 3-5 Broker stakeholder with several SPs and several FPs

Expanded information on the Business Models in T-NOVA will be included in deliverable D8.12, expected by the end of year 2.

4. APPLICATION SCENARIOS

The first part of this section presents two high-level scenarios, to illustrate the use of T-NOVA in two different environments – enterprise and residential.

The section continues with a brief overview of the Virtual Network Functions (VNF) that will be implemented in the context of the T-NOVA project in order to validate the scenarios identified. Along with the motivation for each VNF and an overview of the state-of-the-art, a description and reference to issues and challenges for T-NOVA is presented. Additionally, example usage scenarios for each VNF are briefly discussed. The section addresses the benefits that T-NOVA brings to the deployment and usage of these VNFs.

The third part of this section presents sample application scenarios that involve more than one VNF and are the basis for the future scenarios, to be addressed in later deliverables that provide validation of T-NOVA architecture.

Finally, a discussion on the ETSI NFV Use Cases and their relation to the T-NOVA application scenarios is presented.

4.1. High Level Scenario

This section is aimed at illustrating, through the general description of real world scenarios, how T-NOVA could be used in practice. Two scenarios are provided, targeting different user environments – enterprise and residential.

4.1.1. Enterprise Version

A big corporation called the ABC Company with branches spread across a large geographical area is planning to deploy a virtual private network that will interconnect all the branches with the central offices. Their intention is to create a VPN with specific guarantees for their cross-office traffic, coupled with high bandwidth access to the Internet. In order to decrease their operational costs they would welcome the idea to have some of their security services running in a leased infrastructure. Additionally, they have a requirement to provide a unified communication service for all their personnel, mobile or stationary (located at the offices). Finally, their IT department needs the capability to monitor the infrastructure and have access to granular data on the current traffic/services profile along with data leakage protection.

In order to achieve the requirements above, the IT department administrator accesses the Web interface of the T-NOVA platform. Through this portal, the administrator is capable of providing all the necessary information related to the connectivity service that is required in order to establish the private LAN for interconnecting the branches of the corporation with the central offices. Additionally, the administrator defines the specific functionalities that the network infrastructure would need to support. In fact, the administrator queries the system for a stateful firewall VNF to protect the private LAN from the internet traffic, a Session Border Controller (SBC) for establishing the

unified communication infrastructure and a Deep Packet Inspection (DPI) function in order to collect traffic information and to alert support when abnormal traffic are detected. Specific KPIs related to the performance of the network infrastructure and those of the various VNFs are specified. The T-NOVA platform returns a list of possible solutions for the instantiation of the virtual architecture and the VNFs. If desirable and applicable, a trading/bidding process among multiple actors is initiated with the mediation of a brokerage system, which allows the administrator to get the best price for value for the selected services.

4.1.2. Residential Version

It is weekend time and Alice is looking forward to spending some relaxing time with her friends especially after a rather stressful week at work. As it is raining outside, Alice and her friends decide not to meet in a restaurant and have lunch together, but instead to look for an interesting 3D online game to play together. As her husband (Bob) is finalising a business contract and is planning an important audio/video Conference at 11h in one of the rooms on the first floor, and her son (Anes) is going to watch his favourite education channel through the interactive TV in the living room, she decides to use her tablet and participate in the online game from her bedroom. As a first step, she accesses the T-NOVA Web portal and starts looking for an appropriate virtualised Home Gateway that could fulfil the requirements of the three services to be used by her, her husband and her son. An offer that attracts her attention is based on the service usage and post-payment. This offer also includes a virtualised DPI that could be used to segregate different traffic types and guarantee a better connection to the audio/video conference that her husband needs to handle at 11h, which has priority in case of network saturation. As Bob has been always interested in reducing costs for his company, he simply wants to setup the call on the fly and just pay for the service usage. At 10:45, Bob also accesses the T-NOVA Web portal and starts exploring opportunities for vSBC that he can use to handle his audio/video conference. Since this conference is important and the related content should be kept confidential, he also wanted to have an SBC solution enhanced with some security mechanisms to protect the communication and the content.

4.2. T-NOVA VNFs

Virtualized functions are a fundamental component of T-NOVA. Four VNFs, considered to represent functionalities with widespread use in business or residential environments and high commercial value, are planned to be integrated in the T-NOVA platform. This section provides a general overview of these four VNFs.

4.2.1. Virtual Security Appliance

A Security Appliance (SA) is a device that is used to protect computer networks from unwanted traffic. It can deliver diverse security technologies including firewalling, Deep Packet Inspection, and Intrusion Detection. A Virtual Security Appliance (vSA) is a Security Appliance running in a virtual environment. The type of the security technology is important when it comes to the performance level to be achieved when deploying such technology in a virtual machine.

Nowadays, enterprises wish to integrate devices (including mobile phones) owned by their employees into their enterprise daily activities (email, calendar, documents edition, etc.). This integration clearly raises security concerns, especially in the case that third party software might be installed on these devices. Risks are always present when using devices not belonging to the enterprise, therefore a security appliance is required to protect the assets of the company. If the SA detects suspicious traffic, it will be dropped or redirected to another component for further investigation.

4.2.1.1. Description

In the context of T-NOVA, the vSA is offered to T-NOVA customers that require security, capable of protecting their infrastructure under a 'bring your own device' usage scenario. The vSA can be deployed at the edges of the network, close to the customer premises or at other convenient locations within the virtual network slice that has been provisioned for this customer (see Figure 4-1).

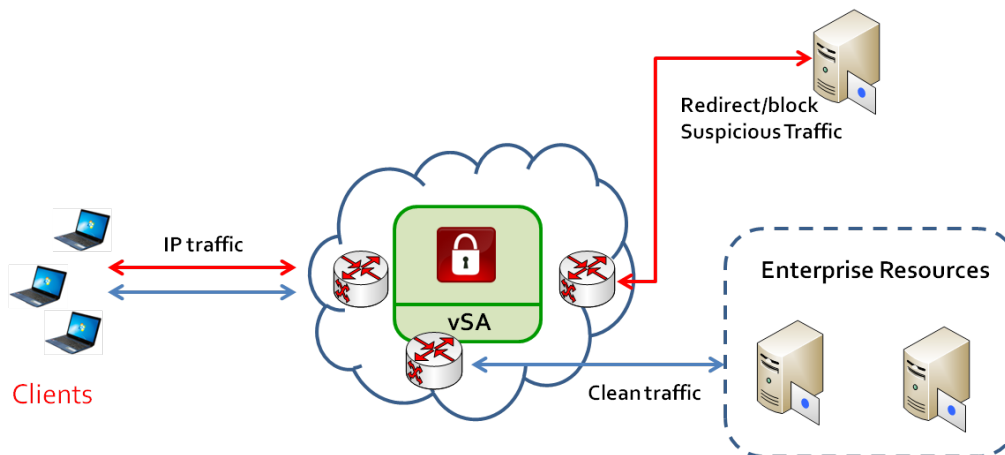


Figure 4-1 Security appliance

The vSA is able to sense potential dangerous or suspicious traffic and responding appropriately to either block it (for example, in case of DDoS attacks) or redirect it to a traffic analysis/forensics virtual device for deeper attack pattern analysis and recognition.

Possible use cases involving SA are:

- Enterprise network protection (from DoS attacks, malware)
- Security support in "Bring Your Own Device" situations.

The virtualisation of Security Appliance requires the optimised network virtualisation at hypervisor level. The packet forwarding performance between the physical network interface and the virtual network interface provided to the virtual machine should be optimised in order to support high traffic volume that requires analysis.

Another trade-off is the performance restrictions. The scaling of the virtualised computing resources and the corresponding workloads pose challenges for the orchestration and virtualisation layer. In this respect, incoming network traffic

variation can create dynamic situations that will challenge the orchestration mechanisms that scale the computing resources in order to preserve appliance performance.

4.2.1.2. Example of Usage: DDoS Attack Mitigation.

Distributed DoS (DDoS) attacks are one of the most significant security problems in the Internet. The problem is well understood, however, its detection is not easy, due to the difficulties in distinguishing between normal and abnormal traffic. Since more and more employees use their own devices to access corporate services, the potential for this form of attack increases if personal employee devices are compromised. The vSA will detect DDoS attacks based on traffic flows information only. This lightweight approach leads to a better performance. The vSA monitors more than one observation point (e.g. vswitch, vrouter), and the statistics collected are to be analysed using expert systems or data mining techniques. If suspicious behaviour is detected, an infection profile is generated and the vSA will initiate a quarantine procedure.

The quarantine procedure is translated into a set of rules that are pushed to the switches in order to disconnect and prevent reconnection of the potentially compromised device to the corporate network.

4.2.1.3. Added-value brought by T-NOVA

The SA was selected for implementation in T-NOVA as the virtualisation of the SA within the context of T-NOVA presents a number of key competitive benefits in comparison to the static hardware appliance approach. These benefits include:

- Elimination of the need for a Customer to accurately dimension the SA capacity and processing needs. The vSA will dynamically scale-up with customer needs. Conversely, most current hardware Security Appliances are considerably under-utilised, leading to increased CAPEX/OPEX as well as wasted energy.
- The vSA can be quickly (even automatically over the Internet) patched with security updates as soon as a new vulnerability/attack technique is discovered or a new mitigation method is designed.
- The virtualised edition of the SA is an excellent candidate for community-driven open source projects, as opposed to "closed" hardware platforms.

4.2.2. Virtualised SBC

A Session Border Controller (SBC) is a device used in multi-media telecommunication providing network interconnection and security services between two IP networks whenever multi-media sessions involve two different IP network domains. In other terms, the SBC is the gateway for multimedia sessions through different networks. It is usually deployed at the border of a service provider network or also in a customer network.

4.2.2.1. Description

An SBC incorporates two separate functions within a single device: the Interconnection Border Control Function (IBCF) for the signalling procedures and the Border Gateway Function (BGF) focused on the user data plane. Signalling procedures are implemented using the Session Initiation Protocol (SIP), while the data or use plane usually adopts Real-time Transport Protocol (RTP) for multimedia content delivery.




Figure 4-2 High level model of an SBC

Possible use cases involving SBC include:

- IP network interconnection
- IPv4-IPv6 gateway
- Business voice services (Enterprise)
- Media transcoding (both audio and video).

The IP network interconnection is the basic service performed by an SBC as a gateway between two different IP networks. The interconnected networks are typically referred as “external” and “internal” or “north” and “south” representing distinct administrative domains. For example, in an enterprise, the “internal” network represents the enterprise private IP internal environment, while the “external” network is the public IP network of the Service Provider or simply the Internet. The SBC has a number of important roles including interconnecting of private and public IP networks, providing topology abstraction and other security related services in addition to NAT (Network Address Translation) required by basic IP interconnection. Moreover, the SBC implements the signalling interworking and media adaptation (transcoding, transrating) functions for effective multimedia communication at the interconnection.

The IPv4-IPv6 gateway scenario is a specific evolution of the previous one where the interconnected networks support also different versions of IP protocol.

Business voice services and Media Transcoding scenarios are related with specific signalling and media adaptation procedures for effective network interconnection. In real deployment it is common that two network uses similar but not identical procedures and protocols. Therefore, the SBC solves the interconnection problem by translating in real-time the various protocols utilised by these networks.

4.2.2.2. Examples of Usage

(a) Interconnecting two sites

The SBC is deployed at the edge of a user network for provisioning multi-media telecommunication interconnection services.

The ABC Company is spread across two sites located in different countries. Each site is connected to the local telco operator. Therefore, each phone call between two sites is charged at international rates.

The management is committed to reduce the cost of telephone calls between sites.

A possible solution leverages on VoIP solution interconnecting the two sites through the public internet using an SBC at the edge of the each site's local company network. Whenever an internal call between the two sites is detected, the direct interconnection will be used instead of contacting the local telco operator.

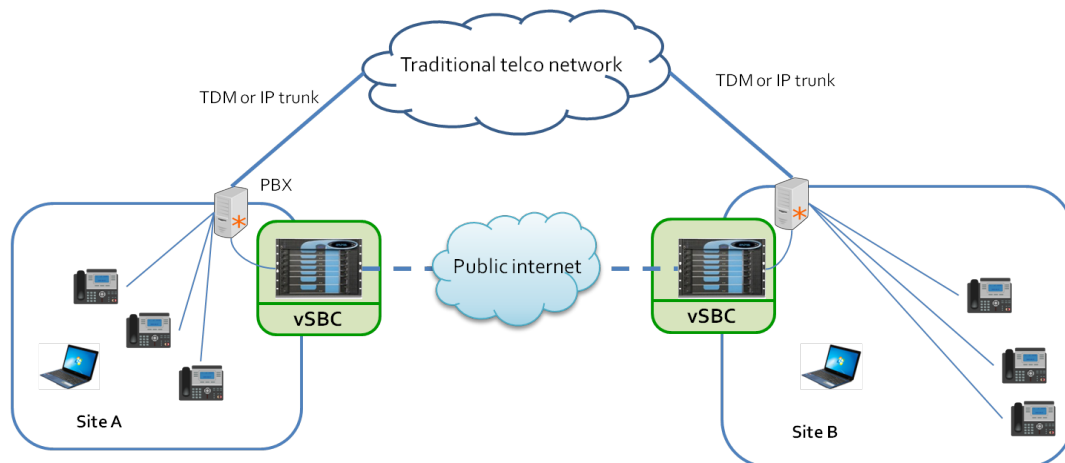


Figure 4-3 SBC interconnecting two sites

(b) Transcoding as a service

The SBC provides media adaptation for video conference.

The company ABC is located at two sites, located in different countries. The sites are already interconnected using a vSBC allowing inter-site telephone calls using direct IP connections through the Internet.

On a day-to-day basis, the normal communications need is for voice calls only. However, periodical video conferences are occasionally organised. Unfortunately, the video communication equipment in the two sites implements incompatible audio/video codecs. Therefore, a real-time transcoding service is required for implementing the video conference.

The solution is to use an instance of a vSBC providing media transcoding service only for the duration of the video conferences. This vSBC will be inserted into the path of the company's inter-site network connection.

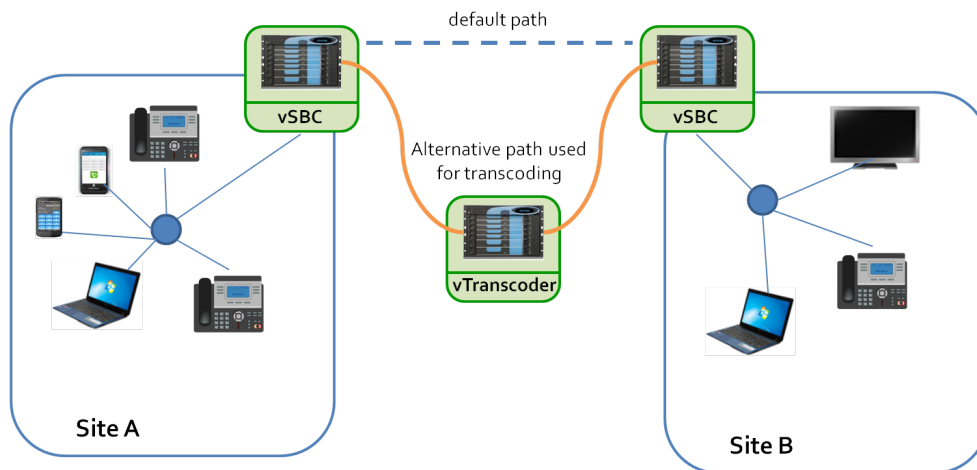


Figure 4-4 Usage of virtualized SBC video transcoding

4.2.2.3. Added value brought by T-NOVA

A Session Border Controller is a very complex and high performing Network Function providing an extensive set of features for effective multimedia communication between different IP network domains. However, each actual deployment requires only a specific subset of SBC features, tailored to specific interconnection needs. Therefore, deploying SBC hardware equipment may result in over-provisioning, in terms of both features and performance.

T-NOVA will simplify the manner in which a customer can purchase the interconnection features he/she exactly needs by selecting the parameters for their specific deployment. Moreover, the customer will be able to change this configuration to meet evolving business needs due to T-NOVA's ability to rescale/reconfigure a VNF such as the vSBC, as required. This flexibility applies to both feature set and performance. Therefore, the initial deployment can be increased or reduced or have its features modified according to actual customer need.

In the usage examples presented so far, an initial requirement for a solution purchased by a T-NOVA customer applies for voice calls only. The customer following the purchase of a SBC identifies that their need for video codec compatibility is only sporadic. This issue is addressed in T-NOVA by adding a virtual SBC configured just with transcoding functionality, which will be used when adding video to a voice call. In this case, T-NOVA detects that the most suitable billing modality would be "pay per use". Therefore, the customer will only incur in expenses when he sporadically uses this transcoding vSBC. This will have a positive economic benefit to the customer when compared to purchasing a new vSBC with all the functionalities and billed on a monthly basis.

In this example, we highlight the benefit of T-NOVA in its ability to offer flexibility in terms of purchasing and configuring virtualised network solutions in the most economical manner for customers, through a preferable billing modality.

4.2.3. Virtualized DPI

Deep Packet Inspection (DPI) is a technology that inspects IP packets at Layer 2 through Layer 7. This includes headers and data protocol structures as well as the actual payload of the message. DPI is used to prevent attacks from viruses and worms at wire line speeds. More specifically, DPI can be effective against buffer overflow attacks, Denial of Service (DoS) attacks, sophisticated intrusions, and a small percentage of worms that fit within a single packet.

A classified packet can be redirected, marked/tagged, blocked, rate limited, and of course reported to a reporting agent in the network. We include both dedicated appliances and embedded Integrated Service Adapters (ISA) for IP security and packet analysis in this definition.

Current market trends reveal that the DPI technology along with policy management frameworks will be deployed as an effective management and network enforcement technology to prioritise traffic, generate new sources of revenue, thwart network disruptions, and meet more stringent regulations on roaming. CSPs in all regions of the world and across residential broadband, mobile, and enterprise business services are implementing policy management systems making use of DPI to enhance services using next generation network technologies.

4.2.3.1. Description

The DPI function automatically recognises application flows. Then, it enables customers to prioritize applications according to their specific requirements, assigning different QoS classes.

For example, within a corporate VPN, flows of real-time applications such as telepresence can be assigned a higher class-of-service than non-real-time traffic such as e-mail or web browsing.

Deep Packet Inspection is a computational intensive activity. One of the requirements for the virtualisation of the DPI functions is the enhancement of the packet processing and handling procedures, in terms of speed and efficiency. The vDPI can elastically grow until a certain upper bound will be reached. It is necessary to combine these goals with an intelligent memory and power consumption system, aimed at a network monitoring system, which will be able to obtain and process a large number of packets quickly, with no additional cost on memory, or complexity.

4.2.3.2. Examples of Usage

(a) Traffic Monitoring

DPI is used for monitoring traffic and generating statistics.

Enterprise users require the ability to monitor their network connections and gather statistical information related to the way network resources are used by their clients/personnel.

In this Use Case, an enterprise customer requests the creation of a specific network slice to interconnect the various office branches and also provide to them access to

the Internet. The customer also requests that he/she should be able to gather information related to the actual information exchanged among the offices and also between the company network and the Internet.

A vDPI is instantiated for monitoring the created network slice.




Figure 4-5 DPI used for Monitoring and Statistics for Enterprise Customers

(b) Traffic Prioritisation

DPI is used for monitoring traffic prioritisation

This scenario deals with another aspect of the DPI, the capability to perform real-time traffic classification for the purposes of traffic prioritisation. In this case, the ingress traffic is steered towards the vDPI and as soon as the inspection of L4-L7 layers concludes, the traffic is identified. This information is used for enforcing the network to prioritise the traffic flow.

A possible technical solution consists of the interaction of a vDPI with an SDN controller in order to handle the traffic according to the policing and prioritisation rules.

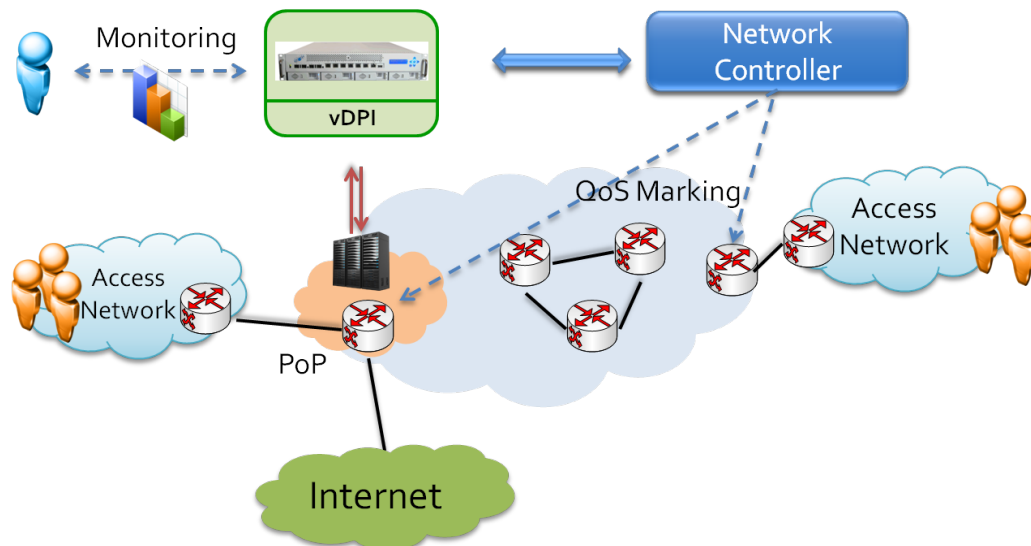


Figure 4-6 DPI used for traffic classification of multimedia streams

4.2.3.3. Added value brought by T-NOVA

DPI enables many functions in the core network, including quality of service (QoS) management through traffic shaping and throttling. Likewise, telcos are creating tiered data plans by using DPI technology to turn the core network into an application aware system. In the context of T-NOVA the DPI benefits can be summed up as:

- Increased flexibility in provisioning and speed of deployment of DPI capabilities within a network. The difference in the deployment speed in contrast to the hardware solutions is significant, in the order of hours versus weeks for hardware based DPI.
- Modularity and on-demand features/functionality during deployment. The capability to provide application awareness coupled with SDN interfacing at the control layer can be used to demonstrate cases where added features (in the form of VNFs) are instantiated on-demand, in order to support traffic patterns or new applications.
- Separation of DPI gateway functions and the management of the DPI within the SDN framework.

4.2.4. Virtualized HG

Physical Home Gateways (HG) are now universally deployed in consumers' house/enterprises. Their main usage is to connect a LAN to a WAN or Internet. They also offer advanced network functionalities like wireless access point, DHCP, NAT, QoS or Firewall. Internet Service Providers (ISP) tend to have a large portfolio of physical devices, depending on the hardware partnership contacts, mergers and acquisitions, device generations, type of Internet connection (xDSL, FTTx, etc...). This high fragmentation results in a variety of costs such as hotlines, supply chain issues, inventory and slow deployment of new functionalities. Moreover, HG software is regularly updated, resulting in unexpected connection outages for customers. For an

ISP, it is not possible to provide High Availability (HA) to their customer with this deployment topology. To circumvent those issues, ISP could use a cloud-based virtual HG (vHG) approach with the benefit of:

- Scaling technically (to a large number of vHGs, possibly in the order of millions) and economically (pay only for deployed vHGs).
- Providing at least the same level of service experienced with current HGs.
- Reducing the fragmentation of deployment configurations.
- Supporting rapid deployment of new functionalities/security updates.

4.2.4.1. Description

A Home Gateway (HG) is a media-centric residential gateway, acting as a logical connection between a SP and Customers. It provides broadband connectivity and multimedia service delivery to a wide set of terminals inside the customer/end-user environment (home or corporate). It allows SPs to supply to Customers advanced context aware multimedia services with the best possible quality.

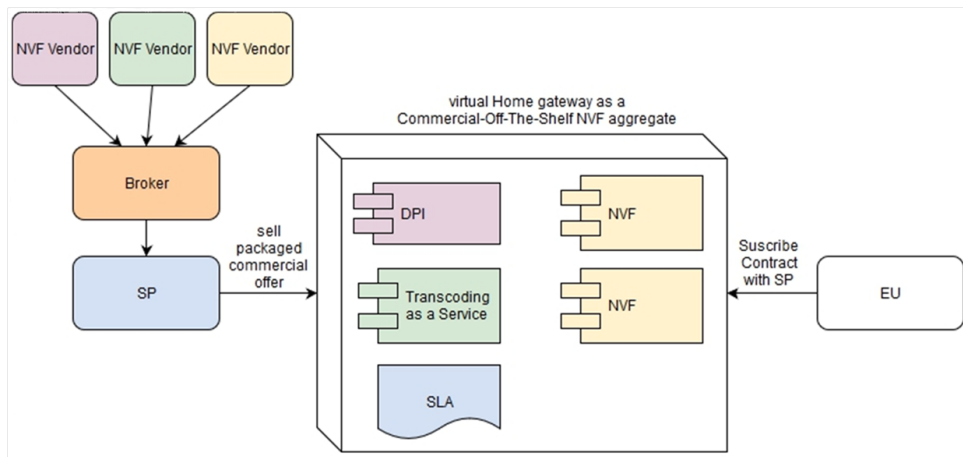


Figure 4-7 Home Gateway

A virtual HG (vHG) is a HG Appliance running in a virtual environment. While increasing the available bandwidth needs between the home/corporate and the network, the vHG provide numerous advantages for providers as presented in the following section.

The main issue for the virtualisation of a HG is transitioning gradually from HG to vHG and scaling massively to very large number of vHG instances. The anticipated signalling volume required to be supported in order to provision and manage vHG instances needs to be considered during design and implementation. This transitioning aspect is also related to the incremental deployment issue, for instance, when customers already have a physical HG that implements part of the functionalities.

4.2.4.2. Example of Usage

(a) Context-aware Service Discovery and Delivery

vHG is used to provide a cost-effective, QoE-driven solution for content delivery

SPs want to supply Customers with advanced context aware multimedia services with the best possible quality and thus increase their revenues. To achieve this, the services need to be first discovered, and then delivered to Customers. Willing to keep up to date with latest business opportunities driven by upcoming technologies, SPs seamlessly update vHG software components and computing power.

For example, a SP decides to implement the promising MPEG DASH (Dynamic Adaptive Streaming over HTTP) streaming technique into a vHG, allowing compatible devices to benefit from improved QoS, and still have legacy devices run in compatibility mode.

The SP R&D uses T-NOVA as a development tool, to benchmark the solution. Results indicate that the new feature requires a 10% increase in compute power in this specific context.

The new feature is deployed during a planned software update on every vHG managed by the SP, additionally the compute power of the vHG is increased by 10% overnight by the T-NOVA Orchestrator.

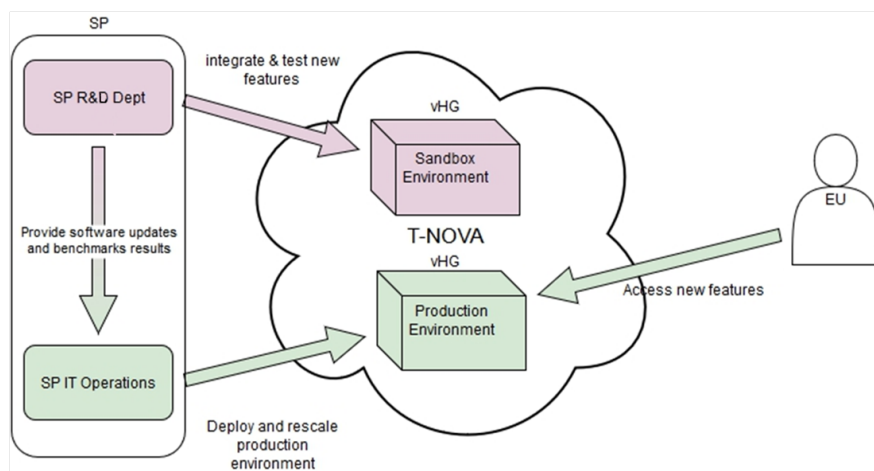


Figure 4-8 HG fast enhancements

The main issue for this scenario is transitioning gradually from HG to vHG and scaling massively to a very large number of vHG instances. This transitioning aspect is related to the incremental deployment issue, for instance, when customers already have a physical HG, which implements part of the NFs. Other issues related to the specific appliances are provided in the corresponding section.

4.2.4.3. Added Value Brought by T-NOVA

A Home Gateway is a necessary but complex device, which provides a set of features for effective multimedia services to home/enterprise Customers. Its virtualisation in

the T-NOVA context, in comparison to the static hardware appliance/deployment, presents a number of key competitive benefits, such as:

- There is no need on behalf of the Customer to care about whether his/her HG is up to date regarding new features/patches from SP and their related processing needs, since the vHG will be automatically updated and dynamically scale-up/down according to their needs. This leads to the vast difference in deployment/maintenance speeds in contrast to hardware based solutions.
- T-NOVA will make it easier for a Customer to purchase third-party HG VNF features (DPI, Transcoding, etc.) they need by selecting them from the T-NOVA Marketplace. The T-NOVA Marketplace and orchestration capabilities will give customers the ability to change the VNF chain according to the services to be consumed.
- The modularity and the on-demand features/functionality deployment by the SP will help to solve the issues related to the incremental deployment of vHGs.
- The virtualised edition of the HG could be also an opportunity for community-driven open source projects, as opposed to “closed” hardware platforms.

4.3. Application Scenarios with more than one VNF

After analysing the four basic virtualised network functions, this section describes concrete application scenarios, in which multiple virtualized network functions are combined to provide advanced services to customers. Like section 4.1, two cases are handled separately – enterprise and residential.

4.3.1. Enterprise Scenario: Attack against the SBC Component

4.3.1.1. High Level Description

The ABC Company is spread in two sites located in different countries. The sites are already interconnected using vSBC allowing inter-site telephone calls using direct IP connections through the Internet. Several people from the ABC Company are located in different locations and need to participate in ad hoc conference call. As they will be using their own devices, the potential risks are high. For instance, a malicious SIP client installed on one of the employee’s mobile device may start flooding the SBC component with SIP messages. The vSA detects this misbehaviour, puts this device in quarantine, and redirects the related generated traffic to a component for further investigation.




Figure 4-9 SA protecting SBC

4.3.1.2. Virtualisation Targets

The virtualisation targets in this scenario are:

- The vSA: comprised of the detection modules and the analysis module (multiple instantiation in various locations at the edges).
- The vSBC: located in the edge of the network close the company sites. This VNF is responsible for providing VoIP services to the company's personnel.

4.3.1.3. Issues / Problems

One of the most interesting issues coming up from this scenario is the calculation of the VNF graph, especially in the case where the topology of the virtual infrastructure is a star topology with traffic from various edges of the network coming towards the SBC. Another issue related to this scenario is the granularity level (IP flows, application layer – signalling protocol -, etc.) when gathering the data that will be analysed for security purposes (for instance, DDoS attacks detection). In addition to that, it is also crucial to ensure an acceptable performance level when analysing the data. It will be also equally important to investigate how a paradigm such as SDN could be used to support discarding the suspected traffic or routing it to another destination for further inspection.

4.3.2. Residential Scenario: vHG with DPI and security appliances

4.3.2.1. High Level Description

The vHG is used to aggregate a specific set of pre-configured VNF and pre-provisioned compute nodes to obtain SPs and EU expectations.

For the general public or basic SME usage, SPs can market vHG packed with the most demanded features including traffic monitoring and security through T-NOVA DPI or Transcoding as a service via the T-NOVA vSBC, as illustrated in Figure 4-10.

From the SP perspective, it alleviates the burden of specific configuration and compute power fine tuning required by aggregating heterogeneous services. It can be marketed as true private cloud solution, in opposition to a more scalable yet more outage-prone public cloud one.

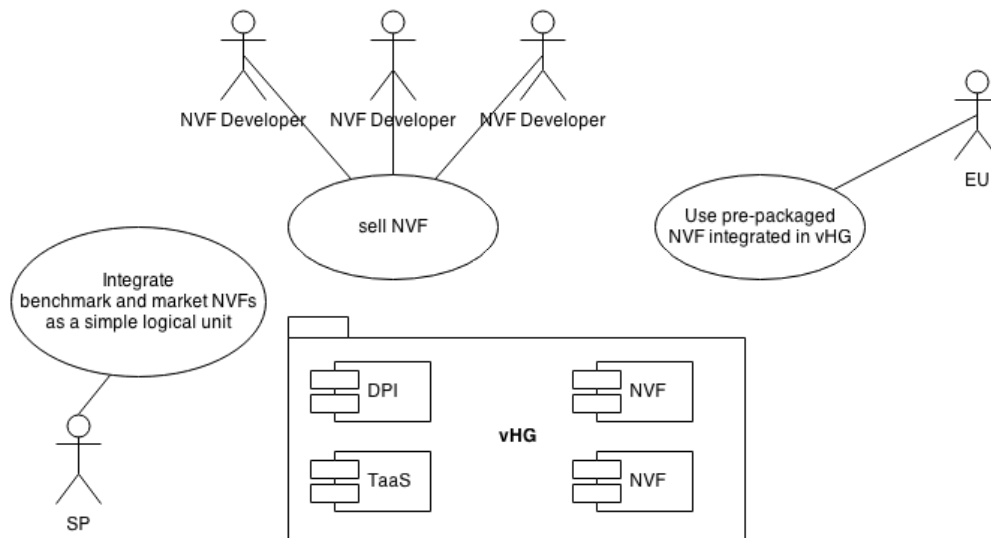


Figure 4-10 vHG with DPI and security appliances

4.3.2.2. Virtualisation Targets

The following virtualisation targets are considered for this scenario:

- The vHG will be the result of the composition of various atomic VNFs. However, the main functionality anticipated is the multimedia service discovery and delivery. Other basic network functions as NAT, DDNS and DHCP will also comprise vHG.
- The vSA will complement the vHG with firewalling and IDS capabilities and at the same time will enable port forwarding exposing LAN machines to Internet.
- The vDPI will be a feature that will enable User Generated Content (UGC) to be detected by the network and allow the proper policing and QoS to be applied on the respective flows.

4.3.2.3. Issues / Problems

The main issues for this scenario are:

- Interface compliance between the appliance specific VNFs and the core VNFs;
- Support of all functionalities which are required for ensuring the secure and QoS-enabled service provision chaining;
- Usability of the appliance specific VNFs in the context where there is a (partial) overlap between their features with those already provided by the core VNFs;
- The composition of all the atomic VNFs into a complete Network Service.

4.4. Relation to ETSI NFV Use Cases

The ETSI NFV Use Cases document (3) describes fields of application, which span the scope of work of the NFV ISG. Both the service models and use cases described in the aforementioned document attempt to provide a commercial and technical context. In this respect, the document is quite relative also for T-NOVA in order to better understand the service models proposed by the NFV ISG and also gather technical requirements and architectures in the future deliverables. In brief, from the

list of Use Cases defined, the following can be considered as more aligned with the T-NOVA scope:

- UC1. Network Function Virtualisation Infrastructure as-a-Service (NFVlaaS). This UC describes the ability of a Service Provider to offer its NFV Infrastructure as a Service (e.g. to other SPs). This enables an additional commercial service offer (in addition to the existing offering of network services offered by supported VNFs) to support and accelerate the deployment of NFV Infrastructure.
- UC2. Virtual Network Function as a Service (VNFaaS). This UC, which is very close to the service model proposed by T-NOVA, offers the VNF as a service provider's application and the enterprise consumer of the service. In this model, the enterprise customer does not need to invest on infrastructure elements or features. Those are provided by the SP, when consuming the VNFaaS. The SP can scale the NFVI resources allocated to the VNF instance in case of increasing demand.
- UC3. Virtual Network Platform as a Service (VNPaaS). In this UC, the SP provides a toolkit of networking and computing infrastructure, as well as potentially some VNFs as a platform for the creation of a virtual network. In comparison to VNFaaS, the VNPaaS provides a larger scale service, typically providing a virtual network rather than a single virtual network function.
- UC4. VNF Forwarding Graphs. This UC demonstrates the complex structure that might need to be supported in the case that VNFs are chained together to form a complete Network Service. The Forwarding Graph defines the sequence of NFs that packet traverses.
- UC7. Virtualisation of the Home Environment. This UC discusses the case where the CPE devices located at home networks (i.e. Residential GWs and Set-top Boxes (STB)) are virtualised. The only prerequisite is the simple, physical connectivity and low cost and maintenance devices at the customer premises.

In the context of T-NOVA, some of the views shared by the ETSI NFV ISG in relation to the aforementioned UCs will be considered. T-NOVA platform provides customers with a connectivity service in addition to NFVI for VNF execution. In this respect, Use Cases 2 and 3 will be addressed. Although T-NOVA does not envisage offering of NFVlaaS by the SP to other SPs, the existence of Marketplace and the associated modules can support the case where more than one SP offer network services through T-NOVA brokering system. In relation to UC4, as already illustrated in the Application Scenarios sub-section, complex scenarios with more than one single atomic VNF are anticipated. In these scenarios, UC4 will be studied and validated. Finally the vHG VNF is actually an absolute match to the UC7 description. In T-NOVA a virtualised multimedia Home Gateway will be implemented and validated.

4.5. Benefits of T-NOVA

T-NOVA VNFs inherits all the benefits of the virtualisation as outlined in section 2. These include, among others:

- Cost reduction;

- Reducing SP's CAPEX by eliminating the requirement to buy a physical device;
- Reducing SP's OPEX by applying pre-configured versions of VNF shaped for the actual customer needs;
- Reducing customer's service cost by charging only for actual use of the service and avoiding large upfront equipment costs;
- Improving energy efficiency;
- Reduced time-to-market;
- Accelerated innovation cycle;
- Improved operational efficiency;
- Easier multi-vendor interoperability.

In addition to virtualisation benefits, T-NOVA introduces the concept of a marketplace where it is possible to choose the most appropriate VNFs aligned with the specific need of a customer:

- Complex services can be implemented by choosing different VNFs that are composed and configured in a flexible manner thanks to T-NOVA system.
- VNFs in T-NOVA marketplace are constantly updated. Therefore, each time a customer needs a VNF service they will have the advantage of having access to the most up to date version.
- When different Function Providers offer similar functions, T-NOVA will offer the customer the best price matching their requirements and SLA needs due to the trading mechanisms among all the function and service providers in T-NOVA.
- T-NOVA helps reducing the overhead when a new service is introduced into an operational network. The new services can be introduced without the need of deploying physical equipment. Moreover, the capabilities of the T-NOVA orchestration framework enables smooth switching to new service configurations.

In this section, we have provided examples of specific application scenarios. These scenarios are presently implemented in networks using physical hardware devices. Finally, we have outlined how these scenarios can be addressed successfully to better meet the needs of customers and service providers by using the T-NOVA system to deliver the benefits of NFV.

5. USE CASES AND REQUIREMENTS

5.1. Methodology

The goal of this chapter is the specification of an initial set of requirements, which is expected to provide guidance to forthcoming stages of the T-NOVA project in the definition of the T-NOVA architecture, its multiple components and ultimately the implementation of the T-NOVA system.

The specification of requirements has followed a use case-driven methodology, as illustrated in Figure 5-1. The requirement specification process included three basic phases, which can be briefly described as follows:

- I. **Business analysis:** the initial phase corresponded to the definition of stakeholders and business roles, as well as the description of a number of illustrative business scenarios and user stories based around the T-NOVA system. The outcome of this work can be found on chapters 3 and 4 of the present document, respectively.
- II. **Use case specification:** this phase included the specification of a number of use cases, describing the interactions between external actors and the system, which are applicable to the business scenarios identified before. These use cases can be found in Section 5.2.
- III. **Requirements specification:** based on the use cases defined in the previous step, an initial set of requirements, addressing different domains, has been identified and specified.

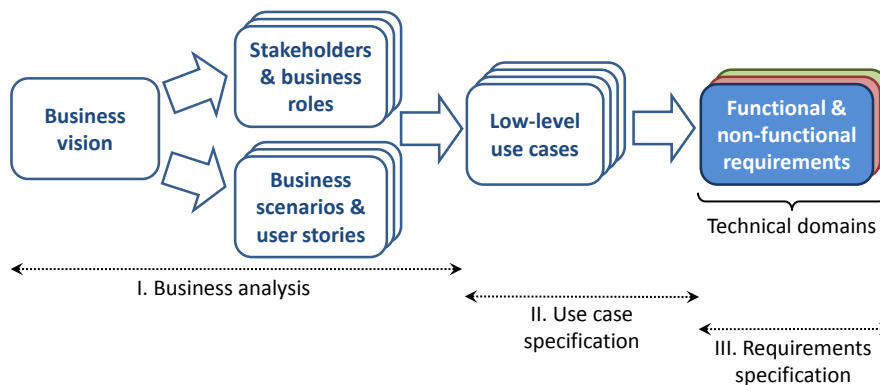


Figure 5-1 Requirement specification methodology outline

At this stage, requirements correspond to a statement of business needs, avoiding as much as possible any implementation bias that may somehow constrain or influence the technical design of the system. It should be noted that these requirements apply to the T-NOVA system as a whole and no assumptions are made about which system components are affected by each requirement, as this will be analysed in subsequent stages of the project.

5.2. Use Cases

Use cases describe the sequence of interactions that take place between the T-NOVA system and the involved stakeholders (which have been identified in section 3), to achieve some outcome of value. Thus, each stakeholder is supposed to have an association with at least one use case.

As mentioned before, the use cases in this section have been specified based on the service lifecycle of VNF services and the associated business scenarios, as defined in the previous chapter. Some aspects of the use cases are expected to be further refined as the project progresses.

It is assumed that business relationships between stakeholders have been established prior to (therefore out of scope of) the execution of the use cases, including the definition of the applicable service parameters and customer profiles.

5.2.1. Basic Use Case diagram

The UML use case diagram depicted on Figure 5-2 (7) represents the T-NOVA roles, use cases and the relationships between them. Ultimately, it can also be seen as a representation of the T-NOVA lifecycle. For a complete description of the T-NOVA stakeholders, please refer to section 3.

Section 5.2.2 provides the detailed description of the use cases that are envisaged to cover the complete T-NOVA service lifecycle.

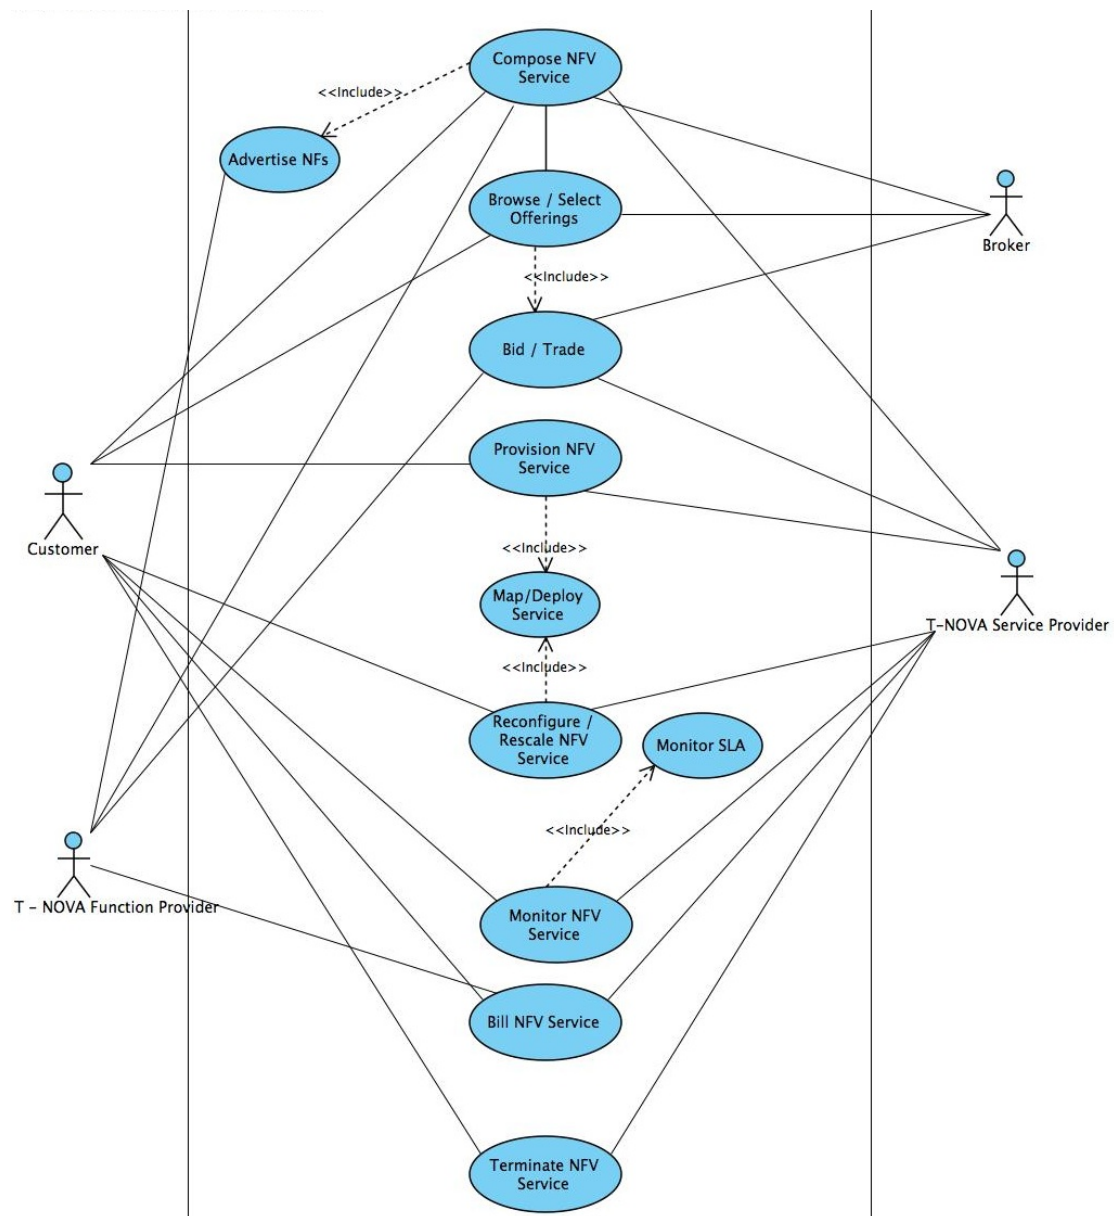


Figure 5-2 T-NOVA Overall Use case diagram

5.2.2. Detailed use case description

This section provides a description of the use cases that are part of the T-NOVA system lifecycle, following a common template, including the following parts:

- Background / Rationale
- Stakeholders involved
- Pre-conditions
- Procedure

5.2.2.1. UC1. Compose NfV services

Background / Rationale

This use case intends to specify the interactions that take place during the VNF composition phase. According to ETSI NFV (5), "VNF composition, is the process

whereby a group of lower-level VNFs is used to define a higher-level VNF". There might be a need for decomposition also. An example would be: The FP offers several smaller functions modules as part of a Security Appliance i.e. Firewall, IPS, IDS. These basic VNFs can be composed in a single Security-oriented VNF, reducing the need for complex service chaining and Network Forwarding Graph calculations.

Stakeholders involved

SP, Broker, FP, (NIP)

Pre-conditions

Two different time frames/phases for this UC can be envisaged. The first one is the off-line composition of VNFs performed by the SP in order to provide bundles or pre-composed VNFs from atomic VNFs as integrated Network Services to their Customers. The second is the real-time dynamic service composition that is triggered by Customer requests through the Dashboard at the Broker. We assume that a Customer has been authenticated at the Broker and exercises the UC1.1 -> UC1.3 sequence outlined below.

Procedure

1. (Offline) The SP advertises atomic VNFs and pre-composed NS.
2. The Broker is authenticated to the system.
3. UC.1.1 - UC1.3 (see following sections).
4. The Broker requests that a new Network Service is composed after a Customer request (as a result of the UC1.3).
5. The Broker submits the desired service components as defined by the Customer (UC1.1).
6. The service is composed, resources are mapped, NFVI-PoPs are identified and the network graph is computed by the SP. The service is decomposed and terminated upon Customer request or following SLA expiration.

5.2.2.2. UC1.1 Browse / select offerings: service + SLA agreement + pricing

Background / Rationale

This use case defines how the customer selects the service among the offerings provided by the SP (service, SLA and pricing), and how the SLA agreement is established among the different involved parties. A contract is established between the Customer and Service Provider, and another between Service Provider and Function Provider (also Infrastructure Provider, if a separate actor), containing (among other things) target service metrics.

Stakeholders involved

Customer, SP, FP, Broker, (NIP, CIP)

Pre-conditions

A T-NOVA customer has authenticated into T-NOVA system and has performed a request.

Procedure

1. The Service Provider, Function Provider (and Infrastructure Provider) create and publish SLA descriptions concerning their own service offerings and customer's obligations.
2. The Customer browses through the different available offerings matching his/her requirements and the SLAs associated to each service as well as the associated pricing. *(In case the service required by the customer is not already offered, a new service composition will take place (UC1) as well as an associated bid/trading mechanism is performed by the Broker (UC1.3))*
3. The Customer selects specific service, SLA and pricing. This can be a bargain-like transaction or simply a combo-list selection of predefined choices.
4. The SP verifies that the Customer profile allows the acquisition of the specific functions corresponding to the selected service.
5. The SP negotiates the appropriate SLAs with the Function Provider (and Infrastructure Provider). This can be done statically in step 1 or dynamically during brokerage.
6. The Customer accepts SLA and pricing (Pricing Module) and other applicable conditions.
7. The SLAs agreed are registered in the system.
8. The system selects the compute/network resources that need to be assigned in order to provide the service meeting this SLA.

5.2.2.3. UC1.2 Advertise NFs**Background / Rationale**

This UC is related to the interactions required for a FP to publish and advertise a VNF.

Actors/Roles

SP, FP

Pre-conditions

There is an offline exchange of authorisation information and certification for each FP, subject to bilateral discussions between the FP and the SP, acceptance of the Terms of Service etc.

Procedure

1. The NF Provider is authenticated to the T-NOVA system.
2. The FP uploads the VNF package, including all the relevant description metadata. The metadata include configuration options, keywords, generic description of the VNF, resources requirements etc.
3. The submitted VNF is certified in order to increase security and integrity of the VNF package. It should be noted that all the steps involved in the VNF deployment should include verification of the certification.
4. The VNF is assigned a unique ID and it is included in the advertised offerings.
5. The FP monitors the status and other statistical data (popularity, voting, comments etc.) of the published VNFs.

5.2.2.4. UC1.3 Bid / trade

Background / Rationale

This use case describes the procedure that is required to perform resource trading among the involved Stakeholders.

Stakeholders involved

SP, FP, Customer, Broker

Pre-conditions

The T-NOVA Customer has been authenticated into the T-NOVA system and set the specifications/requirements of the service that they want to exploit. Either the Broker has returned a list of all potential SPs offering the service, or the Broker has informed the Customer that there is no readily available service matching their needs and thus a new service has to be created. In both cases, the Customer is redirected to a trading process, where SPs bid for their available offerings (e.g. SLA, pricing, etc.), and/or FPs bid for leasing their atomic VNFs.

Procedure

1. The Customer provides their preferences to the Broker, i.e. either the service attributes such as price, duration of usage, billing policy, etc., or the service components and atomic VNFs that the new service may utilise.
2. The Broker advertises these service requirements to the SPs/FPs, receive their initial offerings that match the Customer requirements, and starts a trading process where SPs/FPs are bidding following combinatorial auctions for maximising payoff.
3. The Broker creates the new service portfolio and presents it as a list of new offerings to the Customer.
4. The Customer either selects an offering from this list (step 5), or declines all of them and a new bidding/trading process begins (step 2).
5. The Broker sends the Customer's selection to the corresponding T-NOVA modules for composing the service (UC1.1).
6. The system proceeds with service deployment (UC2).

5.2.2.5. UC2. Provision NFV services

Background / Rationale

The T-NOVA Service Provider instantiates the appropriate infrastructure resources according to the customer request in order to fulfil the SLA.

Stakeholders involved

Customer, SP

Pre-conditions

The T-NOVA Customer has selected the service components and relevant parameters (UC1).

Procedure

1. The SP verifies that the necessary infrastructure resources required to fulfil the customer request are available.
2. Mapping and deployment of the resources are executed (see UC2.1).
3. The SP configures the components of newly created instance(s) of the service.
4. The SP starts the service, verifies that service is up and running and notifies the customer.

5.2.2.6. UC2.1. Map and deploy service

Background / Rationale

The VNFs are mapped into appropriate resources and then provisioned on the NFV infrastructure. The use case may be executed in two different manners – upon a new service request by the customer (UC2), or as a result of a service reconfiguration or rescaling (UC3).

Stakeholders involved

Customer, SP

Pre-conditions

See step 1 of UC2, i.e. the SP verifies that the necessary infrastructure resources required to fulfil the customer request are available.

Procedure

1. The requested service is mapped into specific infrastructure (network and compute) resources, taking into account several applicable objectives (e.g. security, reliability, cost minimization, SLA fulfilment).
2. The virtual network service is established.
3. The images of the requested VNFs are transferred to the appropriate NFVI-PoPs to be deployed.
4. VNFs are instantiated and they are connected to the virtual network.

5.2.2.7. UC3. Reconfigure/Rescale NFV services

Background / Rationale

This UC is focused on the adaptation of the resources allocated to a specific service, optimising resource usage, and/or modification of configuration parameters. The following variants are considered:

UC3.1 Scale-out/ scale-in VNF Service

- Scale-out of the NFV service results in additional VNF instances being added to an existing VNF instance. The new VNF instances require the instantiation of new VMs with compute, network, and/or storage capacity to host the new VNFs.

- Scale-in removes VNF instances and their host VMs that are no longer required. This action releases compute, network and storage resources.

UC3.2 Scale-up/ scale-down VNF Service

- Scale-up results in the compute, network, and/or storage functionality allocated to a specific VNF instance being increased, e.g., replacement of a dual-core with a quad-core processor.
- Scale-down operation results in the compute, network, and/or storage functionality allocated to a VNF service being decreased, e.g. replace XEON E7 processors with XEON E3 processors.

UC3.3 Reconfigure VNF Service

- The configuration/parameters of the service are adjusted.

Scaling may take place either automatically or on-demand:

- *Auto scaling* - if a NFV service is either significantly underutilising or is reaching the maximum utilisation of its allocated compute/storage/network resources, the system will automatically decide to either increase or decrease the allocated resources to the VNF service while maintaining the associated SLA.
- *Scheduled scaling* - Compute resources increased/decreased on a scheduled basis to meet periods of either high and low network traffic.
- *Manual scaling* – The Customer requests increase or decrease in the scale of the VNF service. Alternatively, the SP may manually increase or decrease either the number of VM's or the compute, network and/or storage resources allocated to the existing VMs.

Stakeholders involved

SP, Customer

Pre-conditions

- The Customer has been deployed and is considered active (UC2)
- The VNF service supports auto scale-in/out and scale-up/down

Procedure

Scale Out

1. The Customer issues a scale-out request.
2. The system identifies suitable resources/locations for the additional VNF instances.
3. Resources are allocated and additional VNF instances are instantiated (see UC2.1).
4. The data plane is configured accordingly to connect the new instances.
5. Resource utilisation information and network status is updated.
6. The updated status is communicated to the SP and Customer.

Scale Up

1. The system identifies a VNF service that is over utilising (i.e. reaching saturation point) its allocation of compute/network/storage resources (see UC4).
2. It is confirmed that rescaling is allowed in the SLA.
3. The allocated compute\network\storage resources allocated to the VM(s) hosting the VNF service are increased (see UC2.1).
4. Resource utilisation information and network status is updated.
5. The updated status is communicated to the SP and Customer.

Scale In

1. The Customer issues a scale-in request.
2. The system identifies the VNF instances to be removed.
3. VNF instances are terminated and resources are released.
4. The data plane is configured accordingly to bypass removed instances.
5. Resource utilisation information and network status is updated.
6. The updated status is communicated to the SP and Customer.

Scale Down

1. The system identifies a VNF service that is under utilising its allocation of compute\network\storage resources (see UC4).
2. It is confirmed that rescaling is allowed in the SLA.
3. The allocated compute\network\storage resources allocated to the VM(s) hosting the VNF service are decreased.
4. Resource utilisation information and network status is updated.
5. The updated status is communicated to the SP and Customer.

Reconfig NFV Service

1. The Customer requests a service reconfiguration and submits the new configuration settings.
2. VNFs are re-configured accordingly.
3. Virtual networks are re-configured (topology\capacity adjustment etc.).
4. (If necessary) the service is re-mapped.
5. The Customer and the SP are notified that reconfiguration is successful.

5.2.2.8. UC4. Monitor NFV services

Background / Rationale

The resources consumed by a T-NOVA service as well as its overall status are constantly monitored and monitoring metrics are presented to SP and to the Customer.

The established T-NOVA service is monitored in order to:

- Provide awareness to SP and Customer about service status;
- Provide awareness to SP about infrastructure utilisation;
- Check conformance to SLA;
- Facilitate billing;
- Detect (and possibly prevent) faults and anomalies;

- Trigger reconfiguration/rescaling decisions (UC3).

Stakeholders involved

Customer, SP

Pre-conditions

The service is established and is active (UC2)

Procedure

1. Network monitoring metrics are collected for the connectivity service.
2. VM statistics per VNF instance (CPU, memory, net utilisation etc.) are collected.
3. Monitoring metrics are aggregated to form an integrated picture of the T-NOVA service.
4. Service metrics are presented to the SP and Customer.
5. Service metrics are checked against the SLA (see UC4.1).
6. Anomalies are detected (also possibly forecasted) and alarms are generated.
7. Alarms are presented to SP and Customer.

5.2.2.9. UC4.1 Monitor SLA

Background / Rationale

This use case intends to define the procedures in order to evaluate the agreed SLA between the different parties and to take pertinent actions according to the results.

Stakeholders involved

Customer, SP, FP, (CIP, NIP)

Pre-conditions

SLAs have been agreed between Customer and SP, between SP and FP, (and between SP and NIP/CIP) (UC 1.1).

The service is provided and active (UC 2).

Procedure

1. Service monitoring metrics are collected (UC4).
2. All the terms of the agreed SLA are compared with the metrics provided via monitoring, in order to trace SLA fulfilment and potentially prevent upcoming violations (by warning or requesting more resources – scale-up).
3. SLA results are presented to Customer.
4. SP may also visualise own SLA information, gathering the different providers involved in the service (FP, IP).
5. Billable SLA results are registered for billing, before the bill cycle closes.

5.2.2.10. UC5. Bill NFV services

Background / Rationale

This use case defines the billing procedure for a T-NOVA Customer, and the billing procedure for SP by the FP (and NIP/CIP) based on accounting and SLA fulfilment.

Stakeholders involved

Customer, SP, FP, (NIP, CIP)

Pre-conditions

A Customer is assigned a bill cycle valid for all his subscriptions (a SP is also assigned a bill cycle).

A Customer has requested and selected T-NOVA services, from one or more SPs, possibly with different corresponding pricing conditions and SLAs (UC 1.1).

The service(s) has(ve) been deployed (UC 2).

The bill cycle for T-NOVA Customer is about to close.

Procedure*Billing to Customer:*

1. Accounting information (which includes services subscribed and usages for all services, prices, and billable SLA items) is collected. This is done on a per cycle and per SP basis.
2. The appropriate bills are generated, considering the accounting information.
3. The billing information is presented to the Customer.

Billing to SP:

1. Accounting information (which includes services subscribed and usages for all services and billable SLA items) is collected. This is done on a per cycle and per FP (and IP) basis.
2. The appropriate bills from FP (and IP) to the SP are generated, considering the accounting information.
3. The billing information is presented to the SP.

5.2.2.11. UC6. Terminate NFV services**Background / Rationale**

This use case defines the procedures related to (1) termination of a provisioned NFV service, either by Customer or SP and (2) removal of a VNF from the catalogue of available and advertised services.

Stakeholders involved

Customer, SP, FP

Pre-conditions

- A VNF has been uploaded and published.
- A T-NOVA service has been deployed.

Procedure

Case 1: A Customer is terminating an active/provisioned NFV Service

1. The Customer is authenticated.
2. The Customer browses the active services, chooses the NFV service and requests its termination.

3. The VNF instances are terminated and the service is torn down.
4. All billing/charging activities related to this service are terminated.
5. SP and Customer are informed.
6. NFV related data (monitoring data, billing data, etc.) are archived for further use.

Case 2: A Service Provider is terminating an active/provisioned NFV Service for a specific Customer

1. The SP is authenticated.
2. The SP browses the active services, chooses the NFV service and requests its termination.
3. The VNF instances are terminated and the service is torn down.
4. All billing/charging activities related to this service are terminated.
5. SP and Customer are informed.
6. NFV related data (monitoring data, billing data, etc.) are archived for further use.

Case 3: A Service Provider removes a T-NOVA Service from the service catalogue

1. The SP is authenticated.
2. The SP selects an NFV service from the list of offered services.
3. The SP removes/deactivates the specific service.
4. The NFV service is removed from the list of offered services.
5. The FP is informed.

5.3. Requirements

Following the 3-phase methodology outlined in section 5.1, this section provides the results of the requirements specification phase. The focus is the identification of system-level requirements, driven by use cases. Component-level requirements, following the initial description of system components provided in Section 2, will be specified in future T-NOVA deliverables.

The approach followed in this section is generally in line with the IEEE guidelines for requirements specification (8). By tracing requirements back to its originating use cases, it is possible to understand why every requirement is needed, which stakeholders are involved and which system components are affected.

At this stage, the main focus is the specification of functional requirements to describe the behaviour that the system is expected to exhibit under specific conditions. With few exceptions, non-functional requirements, describing properties or characteristics to be exhibited, or constraints to be respected by the system, have been left to later stages of the project.

The requirements specified stem from the various use cases described in Section 5.2. The ETSI NFV Requirements specification (4) has been taken into account. Requirements address several thematic areas, as follows:

- Management and Orchestration
- Elasticity
- Security

- Resiliency
- Service Continuity
- Operations
- Market / Commercial operability.

When compared against ETSI NFV Requirements (4), T-NOVA has added Market / Commercial operability. Additionally, some areas identified by ETSI NFV have not been considered at this stage and will be addressed later, as necessary:

- Portability / Interoperability,
- Performance,
- Network stability,
- Energy efficiency,
- Migration/co-existence with existing platforms.

Appendix A includes a detailed specification of the T-NOVA system requirements, which are considered to cover the full spectrum of functionalities required by the specified use cases. A total number of 61 requirements have been specified, covering the areas identified above.

The following list summarises the main conclusions:

- **NFV service request.** Customers should be able to express NFV service requests, which will be subsequently submitted to the T-NOVA system. A NFV service specification shall be composed of a set of VNFs advertised by the SPs. For simplicity, a customer should use the Function Store to browse and select among available VNFs. Besides the selection of NFVs, a NFV service request can specify the connectivity of VNFs (i.e., virtual network topology) and any bandwidth or delay requirements for the virtual links.
- **NFV service mapping.** The T-NOVA system should be able to map NFV service requests received from customers to the network, such that all NFV service requirements are met. Specifically, this requires the mapping of virtual network topology to the substrate network, while satisfying any bandwidth and/or delay requirements, as well as the assignment of NFVs to substrate nodes that have sufficient computing and storage resources for packet processing, forwarding and/or caching. In turn, NFV service mapping entails requirements such as the substrate network topology, processing, storage and network resource availability across the network, as well as the computational requirements of the NFVs that should be deployed. NFV service mapping should be optimised based on one or multiple objectives, such as the minimisation of the mapping cost, the maximisation of the provider's revenue or the maximisation of NFV service request acceptance rate.
- **NFV service deployment.** Following NFV service mapping, the assigned computing, network and storage resources should be allocated for the deployment of the NFV service. In addition, the installation of packet forwarding entries is required to ensure that the customer's traffic will traverse the NFVs in the exact order specified in the NFV service request.
- **NFV service scaling.** Existing NFV services should be scaled up or down, upon a customer's request. In the case of up-scaling, this requires the

discovery and allocation of new computing and network resources for the placement of additional VNFs and/or the allocation of more bandwidth in order to accommodate a larger volume of traffic. Conversely, NFV service down-scaling requires releasing allocated resources and possibly the reassignment/reconfiguration of the NFV service in order to achieve resource optimisations.

- **Resource discovery.** NFV service mapping raises the requirement for substrate network topology and resource discovery. Specifically, up-to-date information about the network topology, the bandwidth utilisation as well as the utilisation of the computing and storage resources across the network infrastructure is needed. In addition, the T-NOVA system should have detailed information about the specifications of the VNF hosts, such as the supported VNFs, the number of physical ports, virtualisation technology, etc.
- **Resource isolation.** Resource isolation is a significant requirement for any network services provided on top of shared infrastructures. As such, resources dedicated to collocated NFV services should be isolated from each other. Resource isolation can be achieved using CPU and traffic schedulers in hypervisors.
- **Resource efficiency.** The computational requirements of VNFs can vary significantly, depending on the type of network function (e.g., encryption is significantly more computationally intensive than packet filtering). Therefore, the consolidation of VNFs requires the knowledge of their requirements in order to allocate the required resources and achieve efficiency. This, in turn, raises the requirement for NFV workload profiling.
- **Resource monitoring.** Resource and traffic monitoring is essential in order to achieve resource efficiency and ensure that established SLAs are maintained. Therefore, the T-NOVA system should periodically receive information about the bandwidth utilisation as well as the computing and storage resources utilised by the instantiated VNFs. This information can be used to detect anomalies, resources failures, or severe performance degradation. Such events shall trigger NFV service reconfigurations or reassignments.
- **SLA monitoring.** In order to indicate the status of SLA, the system should be able to compare the service metrics against the SLA requirements. Any violations in SLAs should be promptly reported in order to trigger the necessary actions (e.g., NFV service reassignment).
- **Billing.** NFV services will be offered to different types of customers, such as enterprise networks, service providers, and home network users. Therefore, billing should be tailored to different needs, and, to this end, T-NOVA should support diverse billing models, such as flat rate and pay-as-you-go billing. The latter, in particular, has been proved very successful in cloud computing and is expected to be appealing to many customers.
- **Secure communication and Broker authentication.** NFV service brokerage requires the interaction of SPs with the Broker. To ensure secure communication between these two parties, the T-NOVA system should support mechanisms for Broker authentication and authorisation. Furthermore, the messages exchanged between the NFV Service Providers and the Broker should be encrypted, preventing traffic eavesdropping.

6. CONCLUSIONS

The initial phase of T-NOVA has focused on the specification of use cases and requirements, the definition of basic T-NOVA stakeholders, as well as the description of a number of business-oriented application scenarios and associated value chain, to illustrate how we envisage the deployment of T-NOVA in practice.

One of the objectives of the present deliverable is to establish a common ground on which the remaining WP2 tasks (T2.2 to T2.6), and later the remaining technical WPs (WP3 to WP6), will build their foundations.

This work followed a use case-driven approach, starting with the identification of the participating stakeholders and related business models, the basic use cases, and then evolving to the specification of applicable requirements. The requirements specified in this deliverable should be seen as a first attempt and are expected to be revisited and further refined whenever necessary, in the scope of remaining WP2 tasks and WPs 3 to 6.

From the point of view of the work to be carried out in the future by the several WPs, the specification of application scenarios, use cases and requirements are significant results from this task that will be used as input to subsequent activities of the project. In general, the results included in this document are expected to be utilised in the next steps of the project, in multiple ways:

- Task 2.2 will build the overall system architecture. By defining a common set of technical system requirements, this deliverable lays the foundation for the specification of the T-NOVA architecture and basic service components.
- Tasks 2.3 will concentrate on the specification of the orchestrator platform and its multiple sub-components.
- Task 2.4 will build on the results provided in this document to derive the key technical requirements for the infrastructure and decide on appropriate technological choices.
- Task 2.5 will address the virtualized network functions lifecycle, taking into account, not only the requirements, but also the overall description of the application scenarios included in this deliverable.
- Task 2.6 will build on the initial definition of scenarios and stakeholders provided in this deliverable to develop the concept of T-NOVA marketplace and further elaborate on business models.
- Building on the output from WP2, the remaining T-NOVA technical work packages (namely, WP3 to WP6) will use the requirements defined in this Deliverable as a general guideline for the design and development of new components.

REFERENCES

1. **ETSI NFV ISG.** *ETSI GS NFV 002 v1.1.1 Network Functions Virtualisation (NFV); Architectural Framework.* s.l. : ETSI, 2013.
2. **ETSI.** Network Functions Virtualisation. *ETSI.* [Online] [Cited: 27 5 2014.] <http://www.etsi.org/technologies-clusters/technologies/nfv>.
3. **ETSI NFV ISG.** *ETSI GS NFV 001 v1.1.1 Network Functions Virtualisation; Use Cases.* s.l. : ETSI, 2013.
4. —. *ETSI GS NFV 004 v1.1.1 Network Functions Virtualisation (NFV); Virtualisation Requirements.* s.l. : ETSI, 2013.
5. —. *ETSI GS NFV 003 v1.1.1 Network Functions Virtualisation (NFV); Terminology for Main Concepts in NFV.* s.l. : ETSI, 2013.
6. —. *ETSI GS NFV-PER 002 V1.1.1 Network Functions Virtualisation; Proof of Concepts; Framework.* s.l. : ETSI, 2013.
7. **Booch, G., Rumbaugh, J., and Jacobson, I.** *The UML Reference Manual.* s.l. : Addison-Wesley, 1999.
8. **IEEE.** *IEEE Guide for Developing System Requirements Specifications.* s.l. : ETSI, 1998. IEEE Std 1233.
9. **Bradner, S.** *RFC 2119 - Key words for use in RFCs to Indicate Requirement Levels.* s.l. : IETF, 1997.
10. **TM Forum.** *SLA Management Handbook, Release 3.0.* s.l. : TM Forum, 2011.

LIST OF ACRONYMS

Acronym	Explanation
AAA	Authentication, Authorisation, and Accounting
API	Application Programming Interface
CAPEX	Capital Expenditure
CIP	Cloud Infrastructure Provider
CSP	Communication Service Provider
DASH	Dynamic Adaptive Streaming over HTTP
DDNS	Dynamic DNS
DDoS	Distributed Denial of Service
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
DoS	Denial of Service
DoW	Description of Work
DPI	Deep Packet Inspection
E2E	End-to-End
EU	End User
FP	Function Provider
GW	Gateway
HG	Home Gateway
HTTP	Hypertext Transfer Protocol
IP	Internet Protocol
IP	Infrastructure Provider
ISG	Industry Specification Group
ISP	Internet Service Provider
IT	Information Technology
LAN	Local Area Network
NAT	Network Address Translation
NFaaS	Network Functions-as-a-Service
NFV	Network Functions Virtualisation
NFVI	Network Functions Virtualisation Infrastructure

NFVlaaS	Network Function Virtualisation Infrastructure as-a-Service
NIP	Network Infrastructure Provider
NS	Network Service
OPEX	Operational Expenditure
PaaS	Platform-as-a-Service
PoC	Proof of Concept
QoS	Quality of Service
RTP	Real Time Protocol
SA	Security Appliance
SaaS	Software-as-a-Service
SBC	Session Border Controller
SDN	Software-Defined Networking
SDO	Standards Development Organisation
SI	Service Integrator
SIP	Session Initiation Protocol
SLA	Service Level Agreement
SME	Small Medium Enterprise
SP	Service Provider
UC	Use Case
UML	Unified Modelling Language
vDPI	Virtual Deep Packet Inspection
vHG	Virtual Home Gateway
VM	Virtual Machine
VNF	Virtual Network Function
VNFaaS	Virtual Network Function as a Service
VNPaaS	Virtual Network Platform as a Service
vSA	Virtual Security Appliance
vSBC	Virtual Session Border Controller
WAN	Wide Area Network
WP	Work Package

APPENDIX A. DETAILED REQUIREMENTS SPECIFICATION

To specify requirements, the following template has been used, with the following fields:

Field	Meaning
Req. id	Requirement ID, of the form T_NOVA_xx, in which xx is numbered sequentially, starting from 01.
Use Case	Use case(s) from which the requirement is originated.
Domain	Technical domain to which the requirement belongs, selected out of the list: <ul style="list-style-type: none"> • Management and orchestration • Elasticity • Security • Resiliency • Service continuity • Operations • Market / Commercial operability
Requirement Name	Short requirement name
Requirement Description	Full requirement description. It usually corresponds to a sentence including the word "shall" (for mandatory requirements), ou "should" (for optional requirements).
Justification of Requirement	Rationale behind the requirement
Category	Functional, non-functional

Every requirement has an implicit severity level, which is indicated by the verb used to express it, in accordance to IETF RFC 2119 (8):

- SHALL corresponds to an absolute requirement, something that must be supported by the implementation.
- SHOULD corresponds to a recommended, but optional, requirement – paraphrasing RFC 2119, this means that "there may exist valid reasons in particular circumstances to ignore a particular item, but the full implications must be understood and carefully weighed before choosing a different course".

Generally speaking, the criterion for assessing the severity level of each T-NOVA requirement was basically whether or not that specific requirement is indispensable for the system to deliver its basic function.

Req. id	Use Case	Domain	Requirement Name	Requirement Description	Justification of Requirement	Category
T_NOVA_01	UC1	Security	Authentication and access control	T-NOVA platform SHALL support mechanisms for authentication and authorisation.	Stakeholders interacting with the T-NOVA system should be authorised and authenticated in order for them to browse the Service Offering Database or issue SLA requests	Functional
T_NOVA_02	UC1	Security	Secure communication	T-NOVA customer interfaces SHOULD be secured using encryption.	Encryption should be used, in order to ensure security against eavesdropping.	Functional
T_NOVA_03	UC1	Service Continuity	Network Service De-Composition	The NS SHOULD be decomposed either when the SLA terminates or upon new customer request	The duration of the NS will be specified in the SLA, when the NS is no longer needed the system should de-compose the NS and cancel the SLA. Alternatively the SLA can be terminated by the customer on-demand.	Functional
T_NOVA_04	UC1, UC2, UC3	Operational, Service Continuity, Management & Orchestration	NS Composition	The T-NOVA system SHALL be able to compose a NS from atomic VNF instances available at the NF Store and define the logical topology among the several components.	The creation of a NS from the combination of atomic/simple VNF is important in order to simplify the process provision of NS to the customers and avoid complex path calculations	Functional
T_NOVA_05	UC1.1	Service continuity + Market/commercial operability	Services and SLAs description	The T-NOVA system SHALL be able to allow Service providers and Function Providers to describe their services and conditions (service (+ service description), SLA). One service can be associated with different SLAs; the price cannot be embedded within the SLA.	Services and SLAs description need to be stored in the system to allow the customer to browse through this information	Functional
T_NOVA_06	UC 1.1	Service continuity + Market/commercial operability	Service offerings selection	The T-NOVA system SHALL be able to allow a customer to watch offerings matching their needs, selecting one or more of them. Needs will be mapped to NFV attributes in order to find the most suitable possibilities.	Service, SLA and price information need to be visualised by the customer to allow them to perform a selection	Functional
T_NOVA_07	UC1.1	Service continuity	SLA information store	The T-NOVA system SHALL store SLA agreements among all parties involved (customer - SP, SP- FPs, SP-IPs etc.).	SLA agreements must be stored in order for service monitoring to determine if the SLA has been fulfilled or not	Functional
T_NOVA_08	UC1.1, UC2	Portability / Management and Orchestration	Resource Mapping	The T-NOVA system SHALL be able to map an incoming customer service selection (service + SLA) to specific computational, storage, network infrastructure resources based on specific	Infrastructure resources used to host a specific VNF service and SLA restrictions need to be selected from a pool of infrastructure resources; this selection must be comply with applicable optimization criteria or	Functional

				optimisation criteria or constraints.	constraints	
T_NOVA_09	UC1.2	Security	FP authentication, certification	FPs SHALL be certified and authorised by the system in order to advertise, upload and modify any VNF. The acceptance of an FP is subject of bilateral discussions between the developer and the T-NOVA SP, acceptance of the Terms of Service etc.	Each FP that interacts with the Broker and the Function Store is authorised by the system, by using their private credentials. This is essential in order to control the access to the Broker and the Function Store and increase security.	Functional
T_NOVA_10	UC1.2	Operational	VNF Advertisement	FP SHALL be able to advertise the VNF capabilities in the system.	The FP for each VNF that is uploaded needs to notify the Broker in order that the Service catalogue is updated. This action is called advertisement of the VNF. The advertisement request should contain the name, id, pricing information, requirements and VNF capabilities.	Functional
T_NOVA_11	UC1.2	Operational	VNF Upload	FP SHALL be able to upload the packaged VNF to the Functions Store.	The system should offer a method to the FP for uploading and storing the packaged VNF to the Function Store. When a particular VNF is requested the Orchestrator will instantiate this VNF to the appropriate NFVI-PoP.	Functional
T_NOVA_12	UC1.2	Security/Operational	VNF certification	VNF certification SHOULD be checked.	The submitted VNF is certified by the T-NOVA Function Store in order to increase security and integrity of the VNF package	Functional
T_NOVA_13	UC1.2	Security/Operational	VNF identification	The Function Store SHALL provide a unique identification ID to each certified, advertised VNF.	The VNF id will be the reference name used by the system for monitoring purposes.	Functional
T_NOVA_14	UC1.2	Operational	FP VNF status monitoring	All the VNFs of the same developer SHALL be browsable in the developer dashboard, from where the developer is able to monitor the status and other statistical data (popularity, rating, comments etc.).	This requirement covers the need for supporting the monitoring of each VNF by the FP in terms of availability, popularity, malfunctions and alerting.	Functional
T_NOVA_15	UC1, UC1.3	Service continuity	Service Browsing	The T-NOVA System SHALL be able to browse the available VNFs and services in T-NOVA stores as well as their associated metadata.	The Brokerage must know the available list of Services.	Functional
T_NOVA_16	UC1, UC1.3	Service continuity	Service Trading	The T-NOVA System SHALL be able to carry out trading among several FPs (if there are more than one FPs within T-NOVA providing the same VNFs, or among several SPs).	The Brokerage must carry out the necessary trading in order to provide best suitable results	Functional

T_NOVA_17	UC1, UC1.3	Service continuity	Service Provision	The T-NOVA System SHALL be able to make available the most suitable offerings that the T-NOVA marketplace, can provide to the customer.	The customer must be able to select from a list of Services.	Functional
T_NOVA_18	UC1, UC1.3	Service continuity	Service Auction	The T-NOVA System SHOULD be able to do an auction when the same VNFs or Services or Resources exist from various SP.	The customer must be able to select from a list of Services.	Functional
T_NOVA_19	UC1, UC1.3	Service continuity	Service Composition	The T-NOVA System SHOULD be able to compose a new Service or VNFs in order to match the Customer requirements.	The customer must be able to select from a list of Services.	Functional
T_NOVA_20	UC2, UC3, UC4	Management & Orchestration	Resource monitoring	The T-NOVA system SHALL be able to monitor and collect information about consumption and availability of resources (computational, storage, network) on a real time basis, including the resources consumed by each specific VNF instance.	Monitoring is essential to ensure that the deployment of VNF's onto hosting infrastructure is performed adequately. Monitoring is also r provides essential metrics required by operations such as rescaling, billing, etc.	Functional
T_NOVA_21	UC2	Management & Orchestration	VNF creation	The T-NOVA system SHALL be able to automate the instantiation of VNFs on the infrastructure based on customer request and constraints.	Automation of VNF lifecycle is an essential characteristic of the T-NOVA system	Functional
T_NOVA_22	UC2, UC3	Management & Orchestration	VNF configuration	The system SHALL be able to configure the VNFs running on their host VM's based on a previous customer request for that VNF type or a new configuration specified by the customer	T-NOVA VNFs must be configured following customer request (specific parameters for each VNF)	Functional
T_NOVA_23	UC2	Management & Orchestration	Customer service portal	The customer service portal SHALL provide the means to configure the VNF.	Any parameters required to configure the VNF (e.g. IP prefixes, traffic classes, etc.) must be accessible by the customer through the service portal.	Functional
T_NOVA_24	UC2	Management & Orchestration	VNF start	Once instantiated and configured, a VNF SHALL start its operation on a specific point in time, either on a pre-scheduled basis, or immediately upon customer request.	The customer must be able to determine when a VNF is activated	Functional
T_NOVA_25	UC2	Service continuity	Continuity of basic connectivity service	The activation of the VNF (e.g. insertion in the data path) SHALL have a negligible effect on the basic network connectivity service already in place	The customer must be able to buy and install VNFs without affecting other services already running.	Non-functional
T_NOVA_26	UC2	Management & Orchestration	Topology of VNF components	The T-NOVA system SHALL define the logical topology between the several VNF components.	Connectivity between VNF components must be automated	Functional

T_NOVA_27	UC2, UC3	Management & Orchestration	VNF test	Once the VNF instantiation is completed and the VNF is ready to start, the T-NOVA system SHALL be able to verify that the VNF is operating correctly.	The verification of success of the VNF service creation is needed to provide feedback to the customer and for accounting/billing purposes	Functional
T_NOVA_28	UC2, UC3	Management & Orchestration	Customer notification - VNF starts / fails to start	If the VNF starts correctly, the T-NOVA system SHALL be able to notify the customer about this event. The customer service portal shall provide this information to the customer. If the VNF fails to start correctly, the T-NOVA SHALL be able to notify the customer about this event.	The customer must get feedback about success or failure of his/her service request	Functional
T_NOVA_29	UC2	Management & Orchestration	Accounting notification - VNF starts	If the VNF starts correctly, the T-NOVA system SHALL be able to notify the accounting systems about this event.	For billing purposes, the accounting system has to be notified about the start of the VNF service instance	Functional
T_NOVA_30	UC3, UC4, UC4.1	Management & Orchestration, Operations, Service Continuity	SLA monitoring	The T-NOVA system SHALL be able to compare service metrics with SLA requirements and indicate SLA status (conformance/breach). When the T-NOVA system determines an SLA is in breach it SHALL initiate the applicable action, e.g. rescaling.	SLA management and monitoring is considered essential for the commercial applicability of the T-NOVA system. The T-NOVA system must determine when an SLA is in breach and trigger corrective actions.	Functional
T_NOVA_31	UC3	Management & Orchestration, Elasticity	Customer service portal - Scale In/Out	The T-NOVA system SHALL provide a means for a customer to request either a scale out or scale in of a deployed VNF Service. When the customer requests a VNF scale out or scale in the customer will have the option to reuse a previous configuration or to specify a new configuration.	The T-NOVA system must provide the ability for customers to request additional VNF services or to request the removal of VNF services.	Functional
T_NOVA_32	UC3	Management & Orchestration, Elasticity	Customer service portal - Scale Up/Down	The T-NOVA system SHALL provide the means for a customer to request either a up or scale down the resources allocated to a deployed VNF Service.	The T-NOVA system must provide the ability for customers to request additional resources or the removal of resources from a deployed VNF service.	Functional
T_NOVA_33	UC4	Management & Orchestration	Customer service portal - Reconfig VNF Settings	The T-NOVA system SHALL provide the means for a customer to change the settings of an existing VNF e.g. packet handling rules.	The T-NOVA system must provide the ability for customers to change how their VNF behaves to meet evolving business needs.	Functional

T_NOVA_34	UC3	Management & Orchestration, Elasticity	Customer notification - VNF is removed	A notification SHALL be sent to the customer if the VNF and its host VMs are successfully removed from the T-NOVA system.	The customer must get feedback about the success or failure of their service request	Functional
T_NOVA_35	UC3	Management & Orchestration; Elasticity	Customer notification - VNF rescale	The T-NOVA system SHALL be able to notify the customer when their request to rescale a VNF service has been successfully completed.	The customer must get feedback about the success or failure of their service request	Functional
T_NOVA_36	UC3.1	Management & Orchestration, Elasticity	Customer Scale Out VNF	The T-NOVA system SHALL be able to map an incoming customer service request to scale out an existing VNF service by creating new VMs and deploying VNFs onto the new VMs.	Customers will request increases in VNF services to meet business needs	Functional
T_NOVA_37	UC3.1	Management & Orchestration	Auto Scale Out VNF	The T-NOVA system SHALL be able to automatically scale out a VNF service, based on usage data as well as SLA, by creating new VMs and deploying VNFs onto the new VMs when the SLA associated with the service is breached.	T-NOVA system needs to automatically scale VNF services to meet customer SLA's in an efficient and timely manner.	Functional
T_NOVA_38	UC3.1	Management & Orchestration	Customer Scale In VNF service	The T-NOVA system SHALL be able to map an incoming customer service request to scale in an existing VNF service by releasing resources used by instances of the VNF service as well as hosting VMs, as appropriate.	T-NOVA system must allow customer to request a reduction in their service needs or to completely remove a VNF service as required by their changing business needs.	Functional
T_NOVA_39	UC3.1	Management & Orchestration	Auto Scale In VNF service	The system SHALL be able to scale in an existing VNF service, based on usage data as well as SLA, by deleting VMs due to under-utilization of allocated resources reported by monitoring, if allowed by the SLA.	Ensures that resources are consumed in an efficient manner and SLA specified targets on resource consumption are met.	Functional
T_NOVA_40	UC3.1	Management & Orchestration, Elasticity	Inventory Tracking - VNF Scale Out/In	When a VNF is rescaled, the T-NOVA system SHALL update its inventory of allocated resources.	The T-NOVA System must maintain accurate tracking of resource consumption and details of the services consuming those resources.	Functional
T_NOVA_41	UC3.1	Management & Orchestration, Elasticity	Accounting notification - VNF Scale Out/In	If a VNF is rescaled, the system SHALL be able to account for a change in the allocated VNF service resources for billing purposes.	For billing purposes, the accounting system has to be notified about the additional consumption of resources or reduction in the consumption of resources	Functional

T_NOVA_42	UC3.2	Management & Orchestration, Elasticity	Customer Scale Out VNF	The T-NOVA system SHALL be able to map an incoming customer service request to scale out an existing VNF service by allocating new resources to VMs such as memory and storage or to migrate VNFs from existing VMs to new VMs with higher capability resources such as faster CPUs.	Customers will request increases in VNF services to meet business needs	Functional
T_NOVA_43	UC3.2	Management & Orchestration, Elasticity	Auto Scale Up VNF service	The T-NOVA system SHALL be able to automatically scale up a VNF service, based on usage data as well as SLA, by adding more resources to VMs when the SLA associated with the service is exceeded due to over utilisation of existed allocated resources.	T-NOVA system needs the ability to automatically increase the allocation of resources to maintain customer SLA's	Functional
T_NOVA_44	UC3.2	Management & Orchestration, Elasticity	Customer Scale Down VNF service	The T-NOVA system SHALL be able to map an incoming customer service request to scale down an existing VNF service by decreasing specified amounts of allocated resources from VM's.	Scale down is necessary to ensure the T-NOVA system can meet the changing needs of the customer	Functional
T_NOVA_45	UC3.2, UC4	Management & Orchestration, Elasticity, Resiliency	Auto Scale Down VNF service	The T-NOVA system SHALL be able to automatically scale down an existing VNF service, based on usage data as well as SLA, by reducing the amount of resources allocated to VMs hosting the VNF if allowed under the associated SLA.	Automated scale down is necessary to ensure the T-NOVA system can optimise the efficient use of resources and to maintain cost effective service delivery to customers. Automatic scaling may also be a requirement under the customer SLA.	Functional
T_NOVA_46	UC4	Management and Orchestration, Resiliency	VNF health monitoring	The T-NOVA system SHALL be able to detect anomalies in VNF resource usage and detect possible malfunctions.	To reinforce VNF resilience, which is a critical issue in softwarised infrastructures	Functional
T_NOVA_47	UC4	Management and Orchestration	Monitoring metrics consolidation	The T-NOVA system SHALL be able to aggregate and consolidate all monitoring metrics associated with a service and present the Customer and the SP with an integrated status of the provisioned service.	A consolidated operational picture of the service via the Dashboard is considered a mandatory customer requirement	Functional
T_NOVA_48	UC4, UC5	Operations	Pay-as-you-go billing	The T-NOVA system SHALL be able to exploit usage metrics for pay-as-you-go billing	Pay-as-you-go billing may be considered attractive for some Customers as an option, as opposed to flat-rate	Functional
T_NOVA_49	UC 4.1, UC5	Market / commercial operability	SLA adjustments for billing	The T-NOVA system SHALL store all the information about SLA fulfilment for eventual compensations or penalties.	In order to compensate the customer economically for not achieving the SLA agreed, the billing system must have this information	Functional

T_NOVA_50	UC 4.1	Service continuity	SLA visualisation by customer	The T-NOVA system SHALL be able to allow a customer to visualize SLA fulfilment information when requested.	Customer has to be able to visualize SLA fulfilment information since this is the real service level they are getting and paying for.	Functional
T_NOVA_51	UC5, UC1	Market / commercial operability	Bill cycle	A bill cycle SHOULD be set for each customer that accesses the T-NOVA System.	Billing procedure needs to know when a bill cycle finishes (for non pay- as-you-go services).	Functional
T_NOVA_52	UC 5	Market / commercial operability	Resources usage for billing	The T-NOVA system SHALL store all the information about resources usage by each service for later billing purposes.	Billing procedure needs to know the services that have taken place	Functional
T_NOVA_53	UC 5	Market / commercial operability	Price information for billing	T-NOVA system SHALL store the information about prices agreed by each customer for later billing purposes.	Billing procedure needs to know the price to be applied for service	Functional
T_NOVA_54	UC 5	Market / commercial operability	Bill issuing	The T-NOVA system SHALL issue a bills when the customer's bill cycle finishes or service pay-as-you-go finishes and stores them within the customer profile.	Billing procedure needs to know when a pay-as-you-go service finishes.	Functional
T_NOVA_55	UC5	Market / commercial operability	Bills visualisation by customer	The T-NOVA system SHALL be able to allow a customer to visualise billing information on demand.	Customer has to be able to visualise their bills on demand	Functional
T_NOVA_56	UC6_1 - UC6_4	Management & Orchestration	Service catalogue	The T-NOVA system SHALL be able to provide all services that are active for the authorised customers.	Service catalogue is essential, because it provides to the authorized customers the possibility to identify the service that needs to be terminated	Functional
T_NOVA_57	UC6_1 - UC6_4	Management & Orchestration	Notify/Ack System	The T-NOVA system SHALL notify all relevant participants regarding the removal/deactivation of the service.	Notification System is critical, because it notifies the related users (CSP, NP, End User), when a termination action occurs	Functional
T_NOVA_58	UC6_1 - UC6_3	Management & Orchestration	Release resources, upon removal	The T-NOVA system SHALL be able to release the reserved resources when a removal action is invoked and fulfilled.	Release Resources need to be done, in order to be available for other services and to terminate also the billing actions	Functional
T_NOVA_59	UC6_1 - UC6_3	Management & Orchestration	Archiving service related data	The T-NOVA system SHALL be able to able to archiving service related data for further usage, upon removal of the specific service.	Archiving all related data can be used for further analysis.	Functional