
# Deliverable D4.3:
# Revision of ETICS Architecture and Functional Entities

Deliverable: D4.3

31/01/2012

Version: 1.0
(Final version)


| | |
|---|---|
| Editors: | Ivan Gojmerac and Patrick Zwickl, FTW Telecommunications Research Center Vienna |
| Deliverable nature: | R |
| Dissemination level: (Confidentiality) | PU |
| Contractual Delivery Date: | 31 December 2011 |
| Actual Delivery Date | 31 January 2012 |
| Suggested Readers: | All |
| Total number of pages: | 134 |
| Keywords: | ETICS Overlay Model, Inter-operator QoS, Assured Service Quality (ASQ), ETICS High-level Architecture, ETICS Requirements Assessment |

## ABSTRACT

The substantial challenges in providing interconnection (IC) Quality of Service (QoS) have advanced to an important topic in academia and internetworking practice in the last year. Although the networking industry could potentially rely on mature existing technologies such as the Border Gateway Protocol (BGP), IP eXchange, IP Multimedia Subsystem, (Generalized) Multi-Protocol Label Switching, and the IPsphere B2B framework, none of these approaches has yet been demonstrated to sufficiently support inter-carrier QoS in a way which takes the relevant market requirements into consideration.

In this situation, the present deliverable proposes the ETICS architecture as a self-contained revision of previously introduced ETICS concepts, which aims at enabling a dynamic and automated establishment and management of assured service quality connectivity for the inter-carrier case. With this concept, we pay special attention to the diversity of NSPs in terms of technological heterogeneity, e.g. connection-oriented domains interconnecting with connection-less domains, as well as differences in their business reasoning. Conceptually, this is targeted by the definition of a new kind of network connectivity service, called Assured Service Quality (ASQ) good or ASQ traffic service, which NSPs can offer to their peers, enterprise customers, or Over-The-Top (OTT) service providers – hence making the ETICS services widely accessible in the market.

Methodologically, we start by analysing the shortfalls of the existing technologies, and subsequently we explain the internetworking principles and assumptions which guide and delimit the ETICS solution. This encompasses e.g. the definition of the basic ETICS network services, the introduction of the ETICS "overlay model", as well as considerations about end-to-end ASQ connectivity session services. Based on these concepts, the building blocks of the ETICS system architecture are specified by means of the Unified Modelling Language (UML) accompanied by detailed textual descriptions. Special attention is thereby paid to highlighting the various available deployment scenarios and their options concerning service composition and the publication of offers. Subsequently, the integration of dedicated inter-carrier monitoring solutions and the definition of the Service Enhancement Functional Area (SEFA) complete the description of the ETICS reference architecture.

In a next step, this deliverable provides a preliminary evaluation of the ETICS solution from the perspective of related ETICS project work packages. In particular, recommendations regarding ETICS scalability and performance, as well as the general suitability of the ETICS architecture for various market situations have been extracted. Based on these, the present document provides a set of indications about the relative advantages of the individual architectural options.



**IMPRINT**
Full project title: Economics and Technologies for Inter Carrier Services
Inter-carrier high level technical architecture for end-to-end network services
Document title: Deliverable D4.3: Revision of ETICS Architecture and Functional Entities
Editors: Ivan Gojmerac and Patrick Zwickl, FTW Telecommunications Research Center Vienna
Workpackage Leader: Antonio Cimmino, Alcatel-Lucent Italy
Project Co-ordinator: Nicolas Le Sauze, Alcatel-Lucent Bell Labs France
Technical Project Leader: Richard Douville, Alcatel-Lucent Bell Labs France
This project is co-funded by the European Union through the ICT programme under FP7.

# EXECUTIVE SUMMARY

This deliverable provides a self-contained revision of the ETICS high level architecture, hence rendering the previous architecture document [ETICS-D4.2] obsolete. Therefore, for readers familiar with the previous document, many similarities will be observed. However, the present report significantly goes beyond [ETICS-D4.2] by also discussing the architecture in more detail, particularly thanks to the feedback received from other work packages focusing on business issues, inter-carrier service requirements, more detailed specifications and early simulation and implementation results. An ultimate iteration for evolving the architecture will be performed towards the end of the project to constitute the final ETICS architecture.

The ETICS architecture follows the central objective of enabling the dynamic and automated establishment of assured quality network connectivity services over different network domains owned by different *Network Service Providers* (NSPs). The high level ETICS architecture has to cope with the wide diversity of NSPs, not only in terms of technological heterogeneity (Connection-Oriented (CO) or Connection-Less (CL), forwarding, *DiffServ*, over-provisioning, etc.), but also in terms of business reasoning, i.e. the level of cooperation with other NSPs. The central goal of the ETICS architecture is the creation and management of a new kind of network connectivity service, called Assured Service Quality (ASQ) good or ASQ traffic service, which can be offered by NSPs to their clients, whether they are other NSPs (without the requirement for extending their reachability with service delivery assurance), Enterprise Customers such as Business customers, or Information Services providers (e.g. Communication providers or Over The Top providers). The application of ASQ goods is enabled by the ETICS architecture through supplying the NSPs with a mechanism to find the set of ASQ goods or products that are necessary to satisfy a given customer's connectivity demand or the accumulated demand from a set of customers (e.g. from consumer customers).

Before describing this flexible architecture, the document starts with summarising the known technological solutions such as Border Gateway Protocol, IP eXchange, IP Multimedia Subsystem, Connection-oriented operator Interconnection, and the IPsphere B2B framework (*Section 2*). In each case, we discuss why the existing approaches do not meet the aforementioned ETICS objective, even if ETICS is in part globally benefiting from their broad deployment basis or technological principles. *Section 3* then clarifies the basic ETICS principles, introducing an "overlay model" composed by two interacting and complementary planes: the network service & business plane aiming at the discovery and composition of NSPs' ASQ offers, and the control plane allowing the enforcement of the QoS per-NSP and between NSPs in order to meet technical and business parameters of the ASQ offers. Section 3 also discusses Internet route diversity, which allows finding more paths than the routes propagated by BGP, and it also introduces QoS enforcement principles in CO and CL domains. Subsequently, the main network services envisaged by the ETICS architecture are presented, which are considered having business potential for ETICS providers. An introduction on how ETICS ASQ traffic services can support end-to-end ASQ connectivity session services (microflows) follows, which indicates the potential of how ETICS can support the ASQ needs of end-users and be of relevance in the context of an "ASQ-enabled Internet" supporting both business and consumer customers. Finally,

Section 3 concludes with the discussion of potential policy rules governing the collaboration among participating NSPs, introducing the concepts of ETICS Alliance, Federation & Communities.

*Section 4* represents the core of this deliverable, as it provides an extensive definition of the high level ETICS system architecture. This section presents the processes involved in the business and service plane allowing the cooperation of NSPs. This includes various options for the offer publication and composition, which are needed in order to meet the aforementioned business diversity among NSPs. We see these options as an efficient way to manage the highly complex inter-carrier ecosystem in a single flexible architecture in which these can co-exist. The refined architecture also strongly focuses on the technical modelling of the architectural details by the help of Unified Modelling Language (UML) diagrams, which facilitate the alignment and evolution of preceding results regarding the ETICS implementation and detailed specification. This is furthermore complemented by describing additional components, such as the monitoring mechanism and the *Service Enhancement Functional Area* (SEFA), which provide additional value to the global ETICS architecture. In order to consolidate the ETICS architecture, each option (i.e. how offers are created – push/pre-computed or pull/on-demand – and how offers are composed – fully centralized by a broker, centralized per-NSP or distributed/cascaded) is more detailed technically through the use of sequence diagrams in order to highlight detailed requirements for its implementation.

Finally, *Section 5* provides preliminary feedback from analyses performed in other ETICS work packages on the basis of the initial architecture, in order to start assessing the performance of the various options. Therefore, a quick analysis is done in terms of scalability of the approaches, mainly at the service and business plane which was the focus of deliverable D4.2, showing that both push and pull models bear advantages and weaknesses: while the push model may necessitate handling a large number of (pre-computed) offers, the pull model may impose a higher volume of message exchanges between NSPs, thereby requiring higher computation time. However, no strong technical limitations have yet been observed leading to the exclusion of an option. Similarly, WP3 has provided a business analysis of the applicability of the various ETICS architecture models to a new market of multi-carrier ASQ goods: While a fully centralized model is favoured from a purely economic point of view, practical considerations on the contrary would motivate the use of the per-NSP model for the bootstrapping of the market, as the benefits of a fully centralized approach rely on clear business incentives shared among all NSPs, which is unrealistic in unsettled markets. WP3 and [ETICS-D3.3] thus recommend a stepwise model evolution aligned to market gains in maturity and size, bringing an additional motivation for our flexible ETICS architecture design. Finally, some recommendations have also been extracted from early simulation results (presented in [ETICS-D5.4]) on NSP collaboration models (in particular the sensitivity of possible NSP gains with respect to the way information is exchanged among NSPs), on the benefits of exploring multiple AS-paths to improve the number of served requests (extending therefore current PCE architecture methods), and on the importance of methods to detect SLA violations.

# TABLE OF CONTENTS

# 1. INTRODUCTION

Following deliverable *D4.2 – ETICS Architecture and Functional Entities High Level Design [ETICS-D4.2]*, which has provided a comprehensive picture of the ETICS architecture, this deliverable presents a number of important refinements which have resulted from the continued work both within work package 4 (WP4 – ETICS high-level technical architecture) and from the feedback on D4.2 coming from WP2 (ETICS business and technical requirements), WP3 (ETICS economy & regulation), and WP5 (ETICS detailed technical specifications).

This document thereby aims at serving as a comprehensive, self-contained reference for the ETICS architecture, as it fully incorporates the entire content of D4.2 as well as essential results from related work packages, thus rendering superfluous the need for the reader to additionally consult D4.2.

Previous works within ETICS have identified two main requirements, which need to be addressed by the ETICS architecture as detailed in this report. Firstly, due to the heterogeneity of the infrastructure and market strategies for interconnection services, Network Service Providers (NSPs) are free to use any technology in their network domain (connection-oriented (CO), connection-less (CL), *DiffServ*, over-provisioning, etc.) in order to supply network connectivity services. Secondly, different choices of business models can be made by the NSPs which support such services. The ETICS architecture therefore has to clearly define the inter-carrier interaction mechanisms and parameters that make those kinds of inter-carrier network connectivity services possible, while retaining flexibility in order to enable different business models.

The ETICS architecture is supposed to allow NSPs to find the set of Assured Service Quality (ASQ) goods or products that are necessary to satisfy a given customer connectivity demand. The customer can be an NSP or an end-customer (i.e. an *Information Service Provider*, an enterprise, or even an individual end-user). The flexibility of the architecture, enabling and supporting a range of business models, reflects the context specific requirements of different NSP coordination and collaboration models. As such, the features of the architecture can be adapted to support: (i) carrier oriented network services for large inter-NSP aggregated paths, (ii) business customer oriented connectivity services including cloud connectivity services and VPN services, and (iii) end-user connectivity irrespective of whether explicit session handling is needed or not. These three main categories put different requirements on the inter-NSP interaction processes, for instance regarding the fulfilment process details and response times.

The reader familiar with the previous version of the ETICS architecture document (ETICS D4.2) will also notice that the structure has slightly evolved in D4.3, which accounts for the addition of novel aspects, descriptions, functionalities, and analyses, like e.g. (i) the newly defined *ETICS Overlay Model*, (ii) enhancement of the textual architecture definitions accompanied by detailed UML diagrams, (iii) the definition of the Service Enhancement Functional Area (SEFA), (iv) integration of network monitoring and OAM into the ETICS architecture, (v) analyses of the architecture's scalability, economic aspects and performance, and (vi) the assessment of D2.2 requirements fulfilment by the ETICS architecture.

Accordingly, the rest of this document is structured as follows: *Section 2* provides the reader with an overview of the related work, after which *Section 3* introduces the main assumptions and constraints on

internetworking that delimit the solution space of the ETICS architecture. The definition of the ETICS overlay model thereby represents one of the most important contributions, apart from the definition of ETICS organisational structures like the *ETICS community, ETICS federation* and *ETICS alliance* from a technical point of view. Furthermore, this section also introduces the main types of ETICS network services as traded and managed between the *Network Service Providers* (NSPs), and it also defines the central technical terms which are needed in order to understand the ETICS architecture.

At the very core of this deliverable, *Section 4* introduces the ETICS reference architecture and the service deployment scenarios by providing detailed textual descriptions of the ETICS service composition along with the sequence diagrams necessary for their understanding. Additionally, this section incorporates the presentation of the mentioned *Service Enhancement Functional Area* (SEFA), and it also includes the *session handling* functionalities from D4.2. On a similar note, the descriptions of the network monitoring and OAM have been included into this main section, whereas they had been handled in a dedicated section in the preceding deliverable.

Representing the main link to related ETICS work packages, *Section 5* summarises their feedback on the architecture in crucial areas like scalability, performance, and economic and business aspects. Partially, the feedback described in this section has already been considered and integrated into the relevant architecture descriptions in this document, whereas part of it will be reflected in the final iteration of this series of architecture documents, i.e. D4.4.

Finally, a summary of the main D4.3 conclusions is presented in *Section 6*.

# 2. RELATED WORK

In order to make sure that the ETICS architecture design is based on the current state-of-the-art in IP networking, in this section we review the most relevant related work. We briefly summarise the existing individual solutions, which only partially address the overall Network Service Providers' (NSPs') inter-operator QoS requirements, whereby we highlight the underlying problems that need to be resolved.

Before introducing the high-level architectural choices made by the ETICS project in Section 3, in this section we briefly summarise the technical functionalities missing in the current inter-carrier connectivity approaches, thereby relying on the Sections 2.2 and 2.3 of the previous ETICS deliverable D2.1 [ETICS-D2.1] which have already addressed the issues of current Internet QoS stumbling blocks at a higher level of detail, i.e. in terms of current business models and services, scenarios for the future, high-level requirements and the general question of what the Future Internet may look like.

## 2.1. BORDER GATEWAY PROTOCOL (BGP)

When regarding QoS in an inter-carrier context, a protocol that naturally comes to mind is the Border Gateway Protocol (BGP) as specified in [RFC4271]. Since BGP enables basic inter-carrier layer-3 connectivity in the unique public-IP address space, it can be regarded as the fundamental, underlying service with respect to the additional QoS-enabled traffic transport.

The ubiquitous presence of BGP raises the question of whether it would be possible to extend this QoS-agnostic protocol towards a mechanism for QoS-enabled carrier interconnection. Such extension would offer the advantage of seamless integration of the further required mechanisms with an already existing and deployed protocol, bearing the potential of minimal overhead in terms of additional network management for the novel QoS-enabled connectivity service.

However, there are important stumbling blocks that prevent the integration of inter-carrier QoS connectivity provisioning with BGP. The lack of choice of multiple paths between two remote autonomous systems (ASes) represents one of the most important obstacles. Consequently, each AS chooses and advertises only a single neighbour AS for forwarding traffic towards a specific destination address range, which potentially leads to hard constraints with respect to the QoS capabilities of the given topology, inhibiting the setup of the desired level of inter-carrier quality. An inter-carrier path using a geo-stationary satellite link may serve as a typical example, as it inherently introduces large signal propagation delays, while at the same time other (e.g. fibre optics based) connections between the two communicating ASes may be available.

Beyond the lack of flexibility concerning the choice of inter-carrier network topology, the integration of QoS capabilities with BGP would introduce further scalability issues to this protocol. In other words, BGP currently already suffers from the enormous growth of globally advertised address prefixes (mostly stemming from multi-homed company sites which participate in the BGP routing process), which leads to an explosion of the BGP forwarding table size in backbone routers [BGPMEA]. Additionally, the global advertisement of individual prefix reachability with the current protocol also introduces significant communication overhead, as any information on instabilities of the numerous connections belonging to

multi-homed end-systems is automatically propagated in update messages world-wide, leading to non-negligible overhead in terms of the use of network connections and the computational resources in BGP routers (cf. [RFC4271]).

Nevertheless, more recently two particular efforts seem to have gained some favour at the IETF. The first one concerns the multipath enhancement for BGP that allows an operator to advertise not only the best route, but also alternative ones [IETF-DR-1, IETF-DR-2]. This new feature could be used to generate route diversity, which is a key function in the ETICS framework. The second one aims at allowing BGP to convey *link state* information and more precisely *traffic engineering* information [IETF-DR-3]. Thereby it is important to notice that these two new extensions intend neither to modify the standard BGP protocol nor to modify the routing in the existing Internet. Instead, they mainly intend to provide new functions for the exchange of information between domains that are essential in the context of inter-domain QoS. Even if the ETICS architecture is more based on an *overlay* model (see Section 3), we continue to pay attention to the evolution of these two IETF Internet Drafts.

## 2.2. IP EXCHANGE (IPX)

Having provided an assessment of BGP's suitability for supporting inter-operator QoS-enabled transport, in the following paragraphs we will perform the same type of overview and evaluation for the GSMA standard IP eXchange (IPX) [IR.34], building extensively on our previous description from Section 3.5 of ETICS Deliverable D2.1 [D2.1].

IPX is a further development of the GRX (GPRS Roaming Exchange) which Mobile Network Operators (MNOs) brought to life from the year 2000 on. Back at this time, GRX had been a quite closed network for MNOs with a special purpose (transport of 2.5G and 3G mobile network traffic), but this has changed drastically as the focus of the network has moved from special purpose to more general purpose applications: nowadays, openness is even one of the key factors and envisioned as a business driver of IPX.

Taking into account the three different interconnection models foreseen by IPX, namely IPX Transport, IPX Service Transit and IPX Service Hub, we may conclude that IPX indeed represents a well defined inter-carrier framework for the QoS-enabled exchange of IP traffic, which has led to the fact that international carriers have already started to deploy operational IPX systems [IPX-PR].

As far as the business capabilities of IPX are concerned, the already mentioned three interconnection models do envision different options for the realization of service agreements, financial flows, and technical responsibilities. Building on these capabilities, a variety of inter-carrier business models can be implemented, such that we may conclude that part of the ETICS business requirements could be addressed within the IPX business architecture. However, some aspects of QoS provisioning and the automation thereof seem to be not sufficiently addressed to fulfil ETICS requirements.

It shall be noted that IPX in its own does not mandate specific mechanisms for the provisioning of QoS for end-user sessions and flows, but rather leaves it up to the operators of networks and IPX platforms to make their individual choices (cf. [IR.34]).

Concerning its current status, IPX is increasingly becoming the technology of choice for exchanging high-value IP traffic between mobile network operators. Originally designed only for data traffic of IP (Internet) enabled handsets, IPX has already been further adapted for the transport of VoIP traffic and other types of

more specific services. Even though IPX foresees *DiffServ* as the standard inter-operator QoS mechanism [IR.34], this QoS solution has not yet been adopted by operators, probably due to the difficulties in aligning all operators to a single scheme of QoS classes.

However, as soon as a suitable QoS mechanism is used within individual IPX systems, ETICS solutions could be used to transport the traffic between remote IPX providers with QoS guarantees. In other words, the ETICS system in no way aims at replacing IPX; instead, ETICS rather augments it with an inter-IPX, QoS-enabled connectivity service.

## 2.3. IP MULTIMEDIA SUBSYSTEM (IMS)

Since the uptake of the IP Multimedia Subsystem (IMS) architecture into the mainstream of 3GPP specifications in 2001 (cf. [23.228.Rel5]), the IMS had for many years taken centre stage at the strategic and planning departments of mobile network operators, as the hopes were high that this system would open the door towards the highly desired revenues coming from value added services and mobile handset applications.

This initial vision of IMS becoming one of the main systems in operational networks has however not yet materialized for a number of reasons, the most important of which are: (a) high complexity of the IMS overall architecture which inhibits rapid application development, (b) intransparent roadmap of 3GPP specification development and the set of essential IMS components which has led to functional fragmentation in commercial deployments, (c) the advent of Apple's iPhone and the convincing accompanying ecosystem, followed by other innovative mobile operating systems, offering commercially attractive and consistent application development environments along with matching marketing frameworks.

Whereas the status of the IMS is nowadays still somewhat unclear with most mobile network operators, there does seem to be consensus that the IMS will soon serve as the main platform for carrier-grade voice services (cf. [IR.92]), and that there is still potential to capture at least part of the value added service market using this system.

The perceived potential in terms of value added services is mostly related to the capability of the IMS to assure QoS at the granularity of end-user sessions, opening the way to per application quality assurances. Apart from mobile networks, this situation is also quite similar for the fixed network flavour of the IMS (cf. [TISPAN]), where the IMS could assure the quality in the last mile for premium content traversing the IP data connections of the customers. Support for fixed line networks has been added to IMS in 3GPP release 7 (finalized beginning 2008) which translates not only to opening the network for other (i.e. non-mobile) network operators, but actually also to actively supporting those types of networks in order to improve the range of service coverage and consequently to improve economic attractiveness of the IMS system.

Whereas the IMS does foresee mechanisms for assuring QoS in the access network part, at the same time it lacks similar provisions for inter-carrier sections of the end-to-end path (cf. [23.107], [23.207]). While [23.207] in Annex A.2 does identify a number of scenarios for the realization of end-to-end QoS for IMS-based services, the proposed solutions for the inter-carrier part merely sketch the possibilities for IP QoS differentiation on inter-operator links, and they do not propose concrete mechanisms which would logically associate end-user sessions to the individual classes of traffic at the network edge.

However, when regarded in combination with the previously described IPX, the IMS could very well offer an end-to-end solution for inter-carrier QoS. Whereas a comprehensive summary of the benefits and drawbacks of such an approach is hardly possible to come up with due to a lack of practical experience by the operators even on a pilot project scale, we dare to suggest that while IPX and IMS do represent extensively specified and commercially available solutions, the large overhead they introduce in terms of initial investments, deployment and operations renders this combination a relatively monolithic and quite resource-intensive scheme.

However, there are indications that due to the ever increasing traffic in the access networks, bandwidth problems are emerging also in the core part of the access network, especially on the backhauling trunks. In order to solve this problem, often IMS with its Policy Charging and Rules Function (PCRF) is the solution to this problem by identifying the traffic and shaping it according to its needs. Consequently, the number of IMS enabled networks with PCRF (which is necessary for QoS enforcement) is increasing.

As already explained, IPX could be used to connect IMS enabled access NSPs, but QoS between IPX networks cannot yet be assured. Consequently, instead of IPX, ETICS mechanisms could be used to interconnect IMS access NSPs via ETICS enabled transit NSPs. In this case, the traffic on the edge of the IMS network (whereby the network operator knows exactly about the type of traffic and consequently its QoS requirements) would be put into an ETICS ASQ, and the ETICS system would ensure transport of the traffic with QoS assurances to the remote IMS access network. Within the remote IMS network, IMS-native QoS mechanisms would again take over (provided that the remote IMS network is QoS-enabled), effectively providing for a truly end-to-end QoS-enabled service.

## 2.4. CONNECTION-ORIENTED OPERATOR INTERCONNECTION

Thanks to traffic engineering extensions, Multi-Protocol Label Switching (MPLS) provides network operators with tunnel placement capabilities, thus allowing to select the network QoS parameters per traffic flow. Intra-domain routing protocols (IGPs) are able to carry these parameters and enable automatic path computation engines to select paths across the network with respect to a set of QoS characteristics.

The main limitation of this path computation scheme is related to the fact that, for scalability and confidentiality reasons, traffic engineering parameters are not flooded beyond the boundaries of each IGP area. To address this issue, the Path Computation Element (PCE) function can be used. The PCE architecture provides a communication protocol (PCEP) enabling to request a path calculation from a remote entity, i.e. a PCE. The PCE is typically able to achieve path computation over a larger scope than a usual node, either thanks to a wider knowledge of network information across multiple domains, and/or thanks to its capability to forward the request to one (or several) other PCE(s) which are able to provide the complementary parts of a path. For multi-area networks, a way to use PCE is to implement the function in border routers that are connected to several IGP areas.

However, when it comes to multi-AS and multi-carrier routing, it is unlikely to find an entity in the network which is aware of the necessary information from every sub-network. In this case, a sequence of PCEP requests between all the network domains is to be spanned by the user traffic. The “Backward-Recursive PCE-based Computation Procedure” (BRPC) is a standard specification which uses PCEP to address that problem in case of small interconnections of domains and in case of chains of domains known a priori.

However, there is still a lack in standards on the automation of the calculation of this domain sequence when domains interconnected together are numerous.

In terms of fault detection, usual MPLS mechanisms can be used without change in a multi-AS context, e.g. for loss of signal on equipment interfaces, IGP hellos, Bidirectional Forwarding Detection (BFD), etc. This fault information can then be used to trigger various recovery mechanisms: local repair such as fast re-route or Generalized Multi-Protocol Label Switching (GMPLS) segment protection/restoration, end-to-end protection/restoration (using RSVP-TE Path Error and/or Notify messages to propagate fault detection to head nodes), etc. It shall be noted that there is no known example of current deployment of inter-carrier LSP using signalling protocol messages exchanged between carriers.

Service monitoring usually happens on a per carrier basis, with information relevant to the corresponding AS. The IETF has recently standardized the spatial composition of metrics: [RFC6049] allows combining the traffic metrics from a set of sub-networks or ASes in order to build a complete path metric from the sub-path metrics. What remains to be defined is the way this information is collected among the different ASes and how it is put together; this relationship could be handled by peering or as an agreement part of a consortium of carriers.

## 2.5. SERVICE TO SERVICE PROVIDER AND IPSPHERE

While today's IP transit and peering services are typically not supported by B2B automation, ETICS foresees that the future of inter-carrier (inter-NSP) services will benefit from B2B service management and operations support. An expected result of introducing ASQ inter-carrier services that are reconfigurable is an increase in the number of service instances. Hence, the combination of increased number of service instances and the need to dynamically manage those instances is a driver for automated B2B for service and product management and operations support.

Examples of telco industry operational B2B solutions exist for instance based on the Network Interoperability Consultative Committee[1] (NICC) in the UK, which has developed B2B interoperation standards that enable wholesale oriented automation. This can for instance support management of access line services offered by the access network operator to the retail oriented service providers. So far the focus has been on order fulfilment, trouble ticketing and related support.

The TM Forum appears to be the most relevant international forum for defining, advancing and specifying such B2B frameworks and international de facto interoperability standards. By building on existing TM Forum material, the NICC work and similar efforts in Australia[2] as well as proposals by IPsphere, the TM Forum now is in a good position to evolve the work on general B2B solution frameworks.

The work by the IPsphere Forum and the IPsphere team of TM Forum[3] is an important input to the technical system framework work by ETICS. Although the IPsphere has a wider scope than ETICS in terms of system solution, the overall vision and approach are similar.

The IPsphere Release 1 specification (2007) [IPSPH-R1] is describing high-level requirements, a capability set, and some key interaction patterns enabling advanced and composable inter-provider IP based services.

[1] http://www.niccstandards.org.uk/
[2] http://www.commsalliance.com.au/Documents/national-broadband-network
[3] IPsphere Forum was merged into TM Forum in 2008.

However, the defined interaction patterns and information models where rather rudimentary. No rigorous B2B interaction specification was provided, rather some proprietary specifications have been presented, e.g. as part of demo specifications (TMF IPsphere Catalyst 2009, 2010)[4]. While the IPsphere work and in particular the Catalysts have achieved some well recognized results, the IPsphere framework has not properly taken into account main different network resource levels like large inter-carrier pipes, business connectivity pipes and end-user connectivity or session services. Further, IPsphere has not considered that the management and control of these will require different inter-provider interaction patterns, and it did not raise the issue of how services, processes and information entities at different level can relate to each other. The IPsphere work has so far not addressed monitoring, assurance or charging. However, other TMF work is relevant for these topics. The need for operator specific variants of such a system framework was not sufficiently explored. By modularizing the solution framework is should be possible for the NSP to deploy context specific and targeted instances of the solution framework without having to bear the cost of an all-encompassing solution.

The process of merging IPsphere into TMF (concept harmonization, definitions, etc.) has been long and it is still ongoing. This shows the complexity of the matter and the difficulty of coordinating the many different views from different partners as well as the TMF specific structuring of conceptual and specification work, such as alignment with the different TMF key artefacts (the TMF Frameworx: eTOM[5], Business Process Framework; SID, Shared Information and Data Framework; TAM[6], Application Framework; and the Integration Framework). However, the activities of the IPsphere team have slowed down considerably over the last year. A part of the reason is that IPsphere related topics are considered also in other TMF groups such as the New Services Initiative – B2B working group, in the Service Provider Leadership Council – Wholesale Context working group, the Catalog Management Working Group, and in the SLA Management Team. Moreover, the specification of interfaces as such will eventually be the responsibility of the TIP[7] Team. The decrease in activity in the IPsphere team is perhaps also partly due to lack of operator push, which is again likely to be a result of the well described “deadlock” situation (cf. [ETICS-D2.1]).

ETICS has, through is collaborative efforts across its WPs, defined concepts, information entities and processes that extend and build on the current IPsphere Release 2 architecture [IPSPH-TR158]. The scope of the TMF IPsphere Release 2 Reference Architecture is now more generic in nature compared with the Release 1 work. Note that the Release 2 is a reference architecture setting a frame for further harmonisation with existing TMF artefacts and subsequent specification work. ETICS work and results, on the other hand, are addressing specific services and solutions for interconnect and inter-carrier offerings. Hence, the ETICS concepts and solutions are more specific and targeted than those defined by the previous IPsphere Release 1 document as well as in the ongoing TMF IPsphere work.

ETICS is now in a good position to leverage upon the generic B2B work by IPsphere and TMF, in terms of the existing as well as evolving (by TMF) generic information, process, and interface specifications, and specification methodology. It is important that ETICS in the months to come can in an efficient way progress the ETICS inter-NSP B2B specifications. This will be the target of WP5, based on the updated

[4] http://www.tmforum.org/browse.aspx?linkid=41936&docid=12890
[5] Enhanced Telecom Operations Map.
[6] Telecom Application Map.
[7] TM Forum Integration Program.

architectural and conceptual work in WP4, and the relevant TMF inputs and collaboration leveraging upon the TMF liaison.

However, in the ETICS specification work going forward, we note that there is a risk that trying to achieve the best harmonization with TMF/IPsphere can be very time consuming and prohibit timely and good results from ETICS. Hence, it is a challenge to find a pragmatic and balanced way forward, balancing the time used for learning from TMF and IPsphere and just doing what ETICS believe is a good way forward.

Moreover, it is recognized that the information and data modelling work in particular is both challenging and will set the fundament for the solution development and implementations. Hence, being integral to the specification work, special attention should be put on the ETICS information and data modelling work, taking into consideration different perspectives and harmonization from the general model to the reference point specific models. One goal of this deliverable is to provide a well defined set of concepts and information entity definitions at a sufficient level of detail for WP5 to pick up in the further specification work. Important sources to consider in this respect are the TMF SID and MTOSI models as well as the ITU-T M.3100 [M.3100] and the M.1400-series models [M.1400]. Again the challenge is to keep the model(s) simple and precise while trying to leverage upon the experience and results of previous modelling work and develop this further into the ETICS specific context.

# 3. ETICS INTERNETWORKING PRINCIPLES AND ASSUMPTIONS

The previous section has elaborated on the shortfalls of existing technological options, thereby raising requirements to be addressed by ETICS. In the next subsections, which precede the presentation of the ETICS architecture, the following concepts and definitions of important internetworking principles and assumptions are provided:

- Section 3.1 describes the main issues that the ETICS solutions must solve;
- in Section 3.2, we define how the traffic could be exchanged between carriers, and present the ETICS network services taxonomy (work-in-progress) that will govern the ETICS ASQ IC goods. This includes the definition of the **Point of Interconnect** and the **Traffic Delivery Point** that determine **where** and **how** the NSPs will exchange QoS traffic when internetworking;
- Section 3.3 explains how the network services could be exploited by the network or telco service providers to handle the end-user connectivity sessions (microflows[8]) that will be carried by the ASQ traffic services; and finally
- in Section 3.4, we introduce various kinds of potential ETICS associations of collaborating NSPs. A first attempt at defining a set of **policy rules** is provided that can govern the ETICS framework and accordingly the exchange of QoS traffic between carriers that intend to participate in an ETICS based association.

## 3.1. ETICS SOLUTION PRINCIPLES

We explained in Section 2.1 that BGP has some limitations that the ETICS architecture must overcome. In addition, the design of the architecture faces two non-technical issues:

1. Standardisation bodies, and in particular the Internet Engineering Task Force (IETF), want to avoid the introduction of QoS management in the BGP protocol, mostly for scalability, stability and security reasons.
2. Operators allow only the BGP protocol on the peering points, also mostly for security reasons.

Thus, the ETICS architecture must be built on top of the actual Internet architecture without modifying it. This has to hold in particular with respect to the standard routing mechanisms. This implies that the solution must run as an overlay model, i.e. it has to use an "off-path" signalling[9] approach rather than "on-path" signalling. In addition, the "off-path" approach better adheres to heterogeneous technologies used by the different NSPs, which is beneficial for overcoming the raised constraints on the peering points. It will be easier to achieve consensus on an overlay signalling between NSPs and let them translate this common

[8] The 5-tupple: source and destination IP address, protocol, source and destination port number.
[9] Here, the notion of "signalling" is used in a wide sense. It can involve service management and control.

signalling into their respective underlying technology, rather than aligning all NSPs on the same technology. Even if this is feasible with MPLS-TE, again mostly for security reasons, it will not be possible to use RSVP-TE between the AS Border Router (ASBR) at the peering points to stitch or nest the MPLS tunnels.

This overlay approach has been the guideline of the ETICS architecture design and has been aligned with the definition of the two layers that constitute the ETICS architecture:

1. **(Network) Service and Business plane**: The ETICS network service and SLA management solution is deployed in each NSP domain, communicating over the top of the real topology in order to publish, negotiate, request, monitor and assure the ETICS ASQ traffic services,
2. **Control Plane**: Similarly, enforcement of ASQ traffic services and associated QoS control are performed independently per NSP, At the end, a common end-to-end signalling is used, again off-path, in order to stitch or nest the ASQ part of each NSP to form the composite ETICS end-to-end ASQ traffic service.

### 3.1.1. ROUTE DIVERSITY

The first key point that has to be solved within the ETICS architecture concerns the knowledge of the AS's topology. Indeed, it is not sufficient to only know the best-effort route announced by BGP for a given destination. In order to propose different levels of QoS as well as different price levels, the ETICS system must be aware of all possible alternate routes to reach the given destination. Called "route diversity", this feature must be handled natively by the ETICS architecture in order to have the possibility to compute the optimal route regarding the required QoS and price during the Service Composition step in the SLA life cycle.

Several techniques could be used, like e.g. the add-path feature of BGP. Another option would be to use the Locator/ID Separation Protocol (LISP) [LISP] to advertise some available paths, and consequently the diversity of routes. However, LISP will not announce all possible paths, and the main problem is that no QoS information is attached to a set of multiple paths. Therefore, in ETICS we have selected two different mechanisms to overcome the lack of route diversity:

- **Indirectly** from the SLA offers: In fact, SLA offers contain all pertinent information: QoS & business parameters, network prefixes if they are attached to an edge NSP, border nodes and more, suitable for the service composition to perform SLA abutment and thus provide ASQ offers. The route diversity is obtained by selecting different SLA offers when performing the ASQ composition.
- **Directly** from the network capabilities: In fact, route diversity is obtained from the network capabilities information announced by the different NSPs that participate in the ETICS system. As the network capabilities are conveyed by a link state routing protocol (OSPF-TE or IS-IS-TE), each NSP, at the appropriate AS level, obtains a "global" view of the AS's topology, and thus also the knowledge of all route possibilities for a given destination. Then, the service composition has all pertinent information in hand to select the optimal route regarding the requested SLA.

In both cases, the ETICS solution aims to use an off-path (overlay) signalling (protocol) to exchange pertinent information (SLA offers or network capabilities) to provide the route diversity to the ETICS system

located in each NSP domain. The main advantage in addition to the independency from the underlying network technologies is that no modifications of the standard routing protocol (BGP) are needed.

### 3.1.2. ROUTE ENFORCEMENT

As a consequence of route diversity, the ETICS system must enforce the selected route when it is not identical to the best-effort one during the SLA validation step. Indeed, the routers of each NSP domain compute their respective forwarding table based on the BGP and link state (IS-IS or OSPF) routing protocols, storing best-effort routes only. Once the service composition is done, the ETICS system must enforce the route to be sure that the traffic which belongs to the ASQ follows the selected route, i.e. the one that is described in the ASQ, in order to achieve the appropriate level of QoS. Again, several techniques could be used for that purpose, depending on whether the NSP prefers to choose connection less or connection oriented technologies:

- **Connection Less**: MPLS (without traffic engineering), VPN and GRE tunnel mechanisms could be used to enforce a path for a given set of traffic flows, such that it traverses via the AS Border Router selected during the SLA service composition. While they are simple to deploy, these mechanisms do not provide any support of QoS guarantees. LISP is an alternative solution for enforcing paths towards a given server. In this case, the LISP server must be collocated with the AS Border Router (ASBR) in order to use the correct output of the AS. All these techniques must be used in conjunction with *DiffServ* in order to provide soft (or relative) QoS guarantees. In other words, the over-provisioning management of the network remains vital in this scenario.
- **Connection Oriented**: MPLS Traffic Engineering (MPLS-TE) and in particular *DiffServ*-MPLS-TE technologies provide both route enforcement and QoS support. Of course, it is harder to deploy such techniques; however, they provide great tools for engineering the traffic in a network while enforcing routes with QoS support.

As QoS enforcement is performed per NSP domain (mostly due to the different underlying technologies and also due to the mutual independence of the NSPs), a common off-path (overlay) signalling (protocol) must be used to synchronize the enforcement of the ASQ along the chosen path.

### 3.1.3. QOS ENFORCEMENT

Finally, once the ASQ is enforced in the different NSPs, additional mechanisms must be deployed in order to guarantee the requested level of QoS in each NSP domain. Again, all NSPs are free to select the technologies they prefer to implement QoS. From the ETICS perspective, two kinds of QoS could be supported:

- **Soft or Relative QoS**: This level of QoS is in principal supported if connection-less technologies are deployed in conjunction with *DiffServ*. Connection-oriented technologies could also support relative QoS when no admission control is performed at the entrance of the MPLS-TE tunnels.
- **Hard or Strict QoS**: This level of QoS is only supported when reservation and admission control are performed before the traffic is injected into the network. In addition, connection-oriented technologies facilitate the support of strict QoS in particular with *DiffServ*-MPLS-TE.

The QoS enforcement step could also require off-path (overlay) signalling (protocol) if the configuration is not performed during the route enforcement phase. Indeed, in the particular case of MPLS-TE, both route and QoS are enforced during the same step with RSVP-TE signalling. The only part missing during the configuration phase is to stitch or nest the individual MPLS-TE tunnels (one per NSP domain). But again, this is done during the route enforcement phase (see above).

## 3.2. ETICS TRAFFIC EXCHANGE AND NETWORK SERVICES TAXONOMY

A network service is created by a given NSP based on the knowledge of its infrastructure. Such service therefore initially belongs only to this NSP. The network services of different NSPs need to be stitched or nested together in order to form an ETICS inter-carrier ASQ good. Therefore, it is important to define precisely the boundaries of these network services. In particular, NSPs must agree on **where** and **how** they exchange the traffic that belongs to an ASQ traffic service. To this end, we define boundaries of the NSP's network. We enumerate mainly three levels of boundaries for interconnecting network services:

- The Point of Interconnect (PoI), which is the physical region where NSPs are interconnected,
- The Interconnect Interface (ICI), which is the physical interface by which NSPs interconnect. A PoI may thus be defined by one or several ICIs,
- The Traffic Delivery Point (TDP), which is the attached interface for a given network service on an ICI. An ICI will therefore generally include multiple TDPs.

Once traffic exchanges are localised, NSPs must agree on how they intend to exchange IP packets, and in particular how they will identify the traffic that must follow a given ASQ. To this end, mechanisms for the identification of traffic at interconnection points must be defined.

### 3.2.1. OVERVIEW

In order to enable and facilitate one of the main goals of ETICS, namely that of well-defined inter-NSP B2B interface specifications, it is fundamental to develop and precisely define the set of network services that ETICS supports as well as the main information entities that will be exchanged accordingly on the B2B interfaces to enable efficient product and service offerings and corresponding management, control, and operations support.

The goal of ETICS is to facilitate the provisioning or the supply of inter-carrier services. An inter-carrier service is a network service that spans multiple NSPs. This service results from the concatenation of multiple single-NSP[10] network services. An ETICS customer who desires to obtain an ETICS inter-carrier network service can *request* it from an NSP[11] of the ETICS community. The NSPs within the ETICS community will cooperate according to ETICS inter-NSP interface specifications by exchanging the necessary information in order to find the right composition of single-NSP services that could satisfy the customer request.

[10] Note, in D4.2 the notion of per-NSP was used.
[11] This also includes a virtual NSP or an intermediary acting as a broker.

The goal of this section is to precisely define and present a taxonomy of the ETICS network services, including precise descriptions of the service (end) points and traffic exchange assumptions. The ETICS network service is defined as perceived from the buyer (or requester) and is independent of how the underlying composition is structured or achieved. Note that this taxonomy is however **work-in-progress** and does not cover all potential ETICS network services or service end-points.

An ETICS service is either requested or offered by one of the ETICS actors (cf. Section 4.1.3; Figure 23). Referring to this figure, we therefore place ourselves on one of a set of reference points or "interfaces"[12] between the "supplier" and the "buyer"[13]. While presenting a taxonomy of the services that are exchanged at ETICS reference points and describing some key aspects of each of these types of services, we consider in particular the NSP-to-NSP network services according to E1…E3, and NSP-to-Enterprise Customer (Business Customer or Information SP) services according to E6 and E7. The NSP-to-NSP services according to E4 and E5 are for further study. Moreover, it is assumed that NSP-to-Consumer Customer services can be derived from the NSP-to-Corporate Customer service.

Deliverable D4.2 [ETICS-D4.2] has started a characterization of single-NSP products or services. This characterization mainly focused on the technical (e.g. bandwidth, delay, etc.) and business (e.g. price) parameters that can define a per-NSP product, as well as on the specification of the boundaries that delimit per-NSP products (e.g. traffic delivery points and points of interconnect). This section progresses this early work by elaborating a taxonomy of the ETICS network services with respect to two aspects: who they concern (e.g. a PoI, a host, a region of hosts) and what they can be used for.

As such, the ETICS Network Services Taxonomy attempts to

- define the main types of ETICS network services as traded and managed between the Supplier NSP (potentially also a virtual NSP, e.g. a broker or an intermediary) and
    - the customer NSP (The NSP being part of an ETICS community or alliance), cf. the E1 to E5, and the E1' to E3' reference points;
    - the Enterprise customer (Business customer or Information SP), cf. the E6 and E7 reference points, including the E6' and the E7';
- define the supporting information needed to announce or publish the various types of service or product offers and their associated capabilities. This will include also resource related information representing various types of end-points and service access points. This will include reference to previously defined ETICS SLA templates and SLA management information.

An ETICS network service[14] enables the delivery of traffic from one or a set of points to one or a set of other points. These points represent the boundaries of the ETICS network services. ETICS Network Services are managed by the ETICS framework realized by a set of interacting NSP-specific ETICS systems according to

[12] The notion of reference points and interface is here used interchangeably. However, when speaking of a specific implementation of a reference point one or more technical communication interfaces are applicable. When referring to such specific communication interfaces, based on specific technical protocols, care is taken to clarify this use of the notion of interface.
[13] Note, in some cases the service is requested by a facilitator or client / requester that is not the actual buyer entity.
[14] For short, the notion of ETICS Service or just Service may be use in Section 3. Note that the ETICS framework may provide support services that are not to be confused with ETICS Network Services.

ETICS interface specifications. Whether a subset of the NSPs is part of a multi-NSP alliance obeying to a set of specific business and technical policies is not of relevance to the taxonomy at this stage.

This overview first presents (or recaps) the various types of ETICS network service end-points, after which it provides an overview of NSP-to-NSP services, and the services offered by an NSP to an enterprise customer (a business customer or an Information SP). The last part of this overview considers how consumer customer services (host-to-host) can be enabled by the ETICS network services. More detailed definitions of ETICS services that are within the focus of this deliverable are provided in the subsequent subsections.

We use ETICS network service end-points in this section as the first axis or criteria to perform the service taxonomy. At a high level these points, set of points, or region of points can be one of the following:

- Point of Interconnect (PoI): The PoI identifies a point of interconnection between two NSP networks, i.e. with reference to the E1, E2, and/or E3. A PoI can contain multiple Inter-Carrier Interfaces (ICI) which again can refer to multiple Traffic Delivery Points (TDP). At the network service taxonomy level we may refer to a PoI as the service end-point, while in reality the exact service TDP must be known.
- Point of Enterprise Interconnect (PoEI): The PoEI identifies a point of interconnection between one NSP and an enterprise network (EntNetw). The enterprise network can be a business customer network or an Information SP network. The PoEI will be used as the basis for defining *where* the traffic is delivered (from/to) regardless of *who* would actually "profit" from the traffic delivery service (the "who" is traffic identification space and will be left for future work).
- Host end-point: This point refers to a single host, and can be that of a residential or consumer end-user, a business end-user, an enterprise server host, or a data centre host (e.g. a content server).
- Multipoint: Multipoint is the general term by which we refer to a specific set of given points in terms of either a set of PoIs, a set of PoEIs, or a set of host end-points. In this document, we leave the case of multipoint for future work.
- Region (Destination or Source Region) is a set of host end-points, that is, end-user[15] end-points typically designated by public IP addresses, given in terms of a set of IP prefixes. Note that a specific host within a region may or may not be active at a given time and that ASQ traffic to/from a specific end-point of a region may be controlled by means outside and complementary to the ETICS framework. Hence, a host-to-host inter-carrier network service may not as such be an ETICS network service but rather a network service indirectly enabled by ETICS.

#### 3.2.1.1. ASQ Traffic Charging Principles

The network service types identified in the following are viewed from the perspective of an interconnect point and as delivered by the supplier NSP, either at the carrier PoI or the enterprise interconnect (PoEI). In this respect and in order to remain general we identify two main types of services. Again, note that these are ASQ traffic services at an aggregate level.

[15] When used unqualified an "end-user" can represent any customer (business or consumer) or host including devices and servers.

- Traffic termination (TT): The buyer NSP is sending ASQ traffic to the supplier NSP which is responsible for sending (terminating) the traffic according to the SLA. This is according to the sending party network pays principle [ETICS-D3.2][ETICS-D3.3].
- Traffic origination (TO): The buyer NSP is receiving ASQ traffic from the supplier NSP which is responsible for transporting (originating) the traffic according to the SLA. The constraint on the buyer NSP for the further transport of the ASQ traffic is a topic for further study.

In addition to the TT and TO and in particular for point-to-point pipes such as PoEI to PoEI or PoI to PoI, a dynamic leased line charge for the total bidirectional traffic for the given edge-to-edge or end-to-end service, or its multi- or single-NSP segments, can be considered. This charge can be a traffic dependent charge summing up traffic in both directions, and hence, two associated TT services across a segment can be considered as one service.

In addition, it should also be noted that two different ASQ traffic services for exactly the same destination region and exactly the same QoS traffic class can be charged at different price levels as their respective traffic can be part of different overall value streams. This may be reflected in different service availability requirements for the two ASQ paths.

By considering the different types of “points” mentioned above, a list of the different network service types is presented. First we consider the ETICS network services that are offered and managed over the reference points E1 to E3. Those services are IP or IP/MPLS based services. Similar services are expected in relation to “sub-IP” transport layers over the E4 and E5 reference points. These are considered for further study.

#### 3.2.1.2. NSP-to-NSP (Supplier-to- Buyer/Requestor) Network Services

- PoI to Region (Traffic Termination - TT)
  Buyer NSP is paying for ASQ traffic transported to a given region. Additional constraints will typically apply in order to assure end-to-end ASQ for host-to-host session services as further explained below.
- Region to PoI (Traffic Origination - TO)
  Buyer is paying for receiving ASQ traffic from a given region. Additional constraints will typically apply in order to assure end-to-end ASQ for host-to-host session services as further explained below.
- PoI to PoEI (TT, bidirectional)
  For TT: Buyer NSP is paying for ASQ traffic transported to a given PoEI. Note that several variants of this service apply according to exactly where and how the traffic is terminated.
  For bidirectional: The buyer is paying for two-way traffic across the given segment.
- PoEI to PoI (TO)
  Buyer NSP is paying for ASQ traffic transported from a given PoEI. Note that several variants of this service apply according to exactly where and how the traffic is terminated/originated.
- PoI to PoI (TT, bidirectional)
  For TT: Buyer NSP is paying for ASQ traffic transported to a given PoI. Note that several variants of

this service apply according to exactly where and how the traffic is terminated.
For bidirectional: The buyer is paying for two-way traffic across the given segment.

- PoI to PoI (TO)
  Buyer NSP is paying for ASQ traffic transported from a given PoI. Note that several variants of this service apply according to exactly where and how the traffic is terminated/originated.

For the time being we consider the following network services, which can potentially be offered and managed by means of the ETICS framework, as for further study. Typically, we here speak of an Enterprise server as the host.

- PoI to Host (TT, bidirectional)
  For TT: Buyer NSP is paying for ASQ traffic transported to a given host. Note that several variants of this service apply according to exactly where and how the traffic is terminated.
  For bidirectional: The buyer is paying for two-way traffic across the given segment.

- Host to PoI (TO)
  Buyer NSP is paying for ASQ traffic transported from a given host. Note that several variants of this service apply according to exactly where and how the traffic is terminated/originated.

In addition to the above types of network services we also take into account whether the service is a so-called “Single-NSP” or a “Multi-NSP” network service.

A “**Single-NSP**” network service is a network service where the NSP supplier has responsibility only across his NSP domain. If two neighbour NSPs are bilaterally only offering and buying single-NSP services of each other, the capabilities of the E1 reference point is sufficient for the management of the services. The single-NSP network services are expected to be attractive in a bootstrapping context.

A “**Multi-NSP**” network service is a network service where the NSP supplier has responsibility across his NSP domain and one or more other NSP domains. If two neighbour NSPs are bilaterally offering and buying (cascading) multi-NSP services of each other, the capabilities of the E2 and E3 reference point are needed for the management of the services.

Next we consider the IP or IP/MPLS based ETICS network services that are offered over the E6 or E7 reference points. Similar services, considering connection oriented approach, are expected in relation to “sub-IP” transport layers and still applicable over the E6 and E7 reference points. These are considered for further study.

#### 3.2.1.3. NSP-to-Enterprise[16] (Supplier-to- Buyer/Requestor) Network Services

- PoEI to PoEI (TT, TO, unidirectional or bidirectional)
  Similar options for what is paid for as above.

- PoEI to region (TT)
  Similar options for what is paid for as above.

[16] Enterprise customer is used as a common entity for both the Business Customer and the Information SP (cf. the ETICS reference model).

- Region to PoEI (TO)
  Similar options for what is paid for as above.

For the time being we consider the following network services, which can potentially be offered and managed by means of the ETICS framework, as for further study. Typically, we here speak of an Enterprise server as the host.

- PoEI to host (TT)
  Similar options for what is paid for as above.
- Host to PoEI (TO)
  Similar options for what is paid for as above.

To sum up it is noted that the above set of services enables ETICS to directly manage and support network services at different aggregate levels whether the ASQ traffic service is an inter-carrier pipe between two distant PoIs or an inter-carrier pipe between two enterprise offices (two PoEIs). In addition, the region-based ASQ traffic services can indirectly support session based host-to-host services in particular addressing the consumer customers and their microflows. Moreover, we also observe that for many enterprise services "below" a certain demand or constraint level, it is possible to avoid the need of setting or updating policies on the ASBRs or in the NSP core as the semi-permanent ASQ-enabled NSP core network traffic steering policies are sufficient and there is only a need for configuring or updating policies at the provider-edge (PE) level in a coordinated manner between the edge NSPs.

### 3.2.2. POINT OF INTERCONNECT (POI)

A Point of Interconnect (PoI) represents a "point" in the network where two NSPs interconnect. They delimit the boundaries of each NSP. PoIs often reflect geographical locations and can be associated with a Point of Presence (PoP) at which two NSPs interconnect. As illustrated in FIGURE 1, two NSPs can be interconnected with one or more PoIs. In the case of the Internet, this happens when two ASes are connected to each other at more than one geographical location or PoP.

In addition, at a given PoP there can be several NSPs present and hence there can be several PoIs associated with a PoP. For example, in FIGURE 2, a single Autonomous System Border Router (ASBR) both serves for a connection between NSP A – NSP B, and for the connection NSP B – NSP C: this ASBR is thus part of two PoIs.

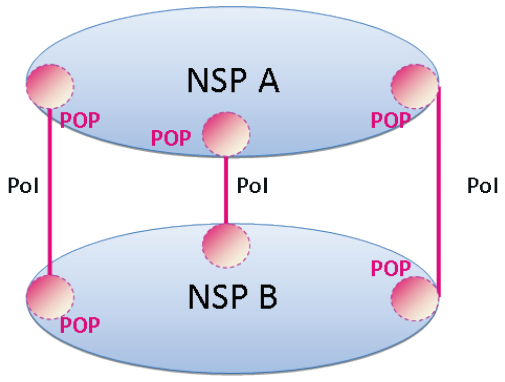


FIGURE 1: POINTS OF INTERCONNECT

Points of Interconnect must be identified during the Service composition phase. Indeed, when the service composition process will try to abut SLA offers, it must be aware about how to do this step. For that purpose, we intend to use the PoI as the identifier within the SLA offers to determine if two SLA offers could be combined or not in order to form the end-to-end SLA. There are two ways to identify the Point of Interconnect:

- By **location name**: Exchange point or facility name are used to name the Internet peering exchange point. PoI could use the same wording, especially if they intend to use the same peering point as the one used by the Internet. If two SLA offers, coming from two NSPs, used the same name as the PoI to localise their respective AS Border Router, then these two SLA offers can be combined. If this convention is used, it is mandatory to additionally maintain a corresponding table between the PoI name and the IP address of the AS Border Router (ASBR) loopback interface. This is needed for the ASQ enforcement step which is used to configure the ASBR. For that purpose, such IP address must also be part of the SLA offers in order for the ETICS system to be able to enforce the ASQ.
- By **IP address**: In fact, AS Border Routers use a dedicated loopback IP address to advertise themselves, at least within the BGP protocol. These IP addresses of all the AS Border Routers located in the same PoI could be used to name the PoI itself. While this would bear the advantage of directly embedding the IP address of the AS Border Router for the ASQ enforcement, at the same time it would increase the complexity of performing SLA offer abutment. In fact, this solution is more accurate when exchanging network capabilities information. Indeed, the IP address is automatically announced by the link state protocol, and thus made available to the service composition. Then, when using a Path Computation Element (PCE) for the connection-oriented technology, the PoI will be provided to the PCE in the form of the Include Route Object (IRO) parameter in the PCE request, in order to explicitly designate the AS Border Router, i.e. the PoI, which must be followed by that path during the BRPC computation.

Both cases will allow the ETICS system to select and impose the PoI that must be used during the SLA service composition. Thus, the ASQ will go through the PoI that corresponds to the optimal route regarding the original SLA request.

### 3.2.3. INTERCONNECT INTERFACE (ICI):

A PoI that connects two NSPs might concern more than a unique couple of ASBRs. We therefore introduce Interconnect Interfaces (ICI) as the physical links connecting peering ASBR network interface cards.

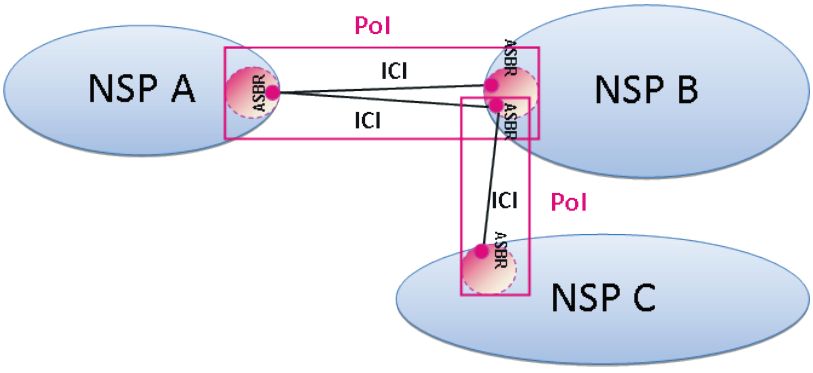


FIGURE 2: INTERCONNECT INTERFACES

FIGURE 2 illustrates the terms: NSPs A and B are interconnected by a PoI, and NSPs B and C by another PoI. The PoI between NSPs B and C only involves one ASBR per domain, and this includes a single ICI. On the contrary, NSPs A and B can be interconnected thanks to two ASBRs on B side in the same PoP, so two ICIs can be defined.

### 3.2.4. TRAFFIC DELIVERY POINTS (TDPS):

Once an inter carrier service has been instantiated, the traffic is forwarded between NSPs. Each NSP is responsible for the quality it has promised for the product it sold. It is important therefore to precisely define the boundaries that define the areas where the responsibility of each NSP applies. In particular, NSPs must agree on where, which amount and which part of traffic is concerned by a given SLA contract. NSPs need therefore to agree on reference points where measurements can be performed in order to check the compliance of the contracts (SLAs and SLSs).

In order to facilitate the service selection and composition, the SLA negotiation and management, as well as the SLS characterization, ETICS has defined some particular points in the network topology to serve as "Traffic Delivery" boundaries.

#### 3.2.4.1. Traffic Delivery Point Definition

Located at the Point of Interconnect (PoI) on one identified ICI or at an interface of a provider edge node, the "Traffic Delivery Point" (TDP)[17] is associated with one ASQ path product instance, i.e. with one SLS instance of that product, and describes several fundamental characteristics of the traffic for that SLS instance belonging to a given SLA contract:

- means that enable the unique identification of the traffic,
- quality performance commitment applied to this traffic,
- ASQ traffic ingress and the egress points for the SLS instance, and
- any associated transit traffic delivery point (if applicable), possibly considering upstream and/or downstream delivery points.

An SLS therefore applies at a TDP and describes the QoS objectives in the transport of the traffic that the NSP must respect between this TDP (ingress) and the next TDP(s) (egress). Indeed, an NSP accepts traffic at a given TDP and transports it to the next TDP. NSPs are free to choose any QoS mechanism and transport technology they prefer; however, they must respect their SLS parameters commitments between two TDPs (ingress and egress). The list of TDPs for a given end-to-end end-customer SLS (that we call *EC-SLS*) defines a chain of ASs/NSPs between the source and the destination. The end-to-end QoS objectives are thus divided into per-NSP/AS SLSs (that we call *N-SLS*), each of them defined between TDPs of the different ASs/NSPs along the chain and representing the individual contracts of each NSP.

Because an NSP can manage several ASes, TDPs between two AS belonging to the same NSP could be seen as internal (from the NSP point of view) as opposed to TDPs between two ASes belonging to two different NSPs. However, from the ETICS perspective, no distinction will be made, as the purpose of service composition is to at least compute the AS chain (at least through identification of PoIs) regardless of which

[17] The TDP taxonomy first introduced in D4.2 is currently still under revision.

NSP manage the ASes. All other internal NSP reference points are not handled in this definition and NSPs can manage their internal reference points in any way they choose.

#### 3.2.4.2. TDPs Taxonomy

There are two main categories of TDPs: Transit TDP and Termination TDP. In ETICS, for each service instance we clearly distinguish between the source and the destination points. The source refers to where the traffic is taken in charge of by the first NSP (if the service relies on several NSPs/ASs). The destination denotes where the traffic is delivered by the last NSP (if the service relies on several NSPs/ASs).

At the source, the customer must provide its traffic in conformity to the signed SLA contract if it wants it to be carried to the destination with the negotiated QoS. At the destination, the destination TDP is the point where the end-to-end measurement could take place to verify the SLA contract and in particular the end-to-end QoS commitment.

The Transit TDPs are only located between NSPs.

Termination TDPs are differentiated according to whether they are at the source (ingress) or at the destination (egress) but are also according to the nature of the source or destination:

- TDSP – Traffic Delivery Source Point: TDP that is the source of traffic delivery. Source is here a single point.
- TDDP – Traffic Delivery Destination Point: TDP that is the destination of traffic delivery. Destination is here a single point.

However, in some cases, the source or the destination of the traffic may not be a single point but a group of points, therefore called a region. In these cases, we cannot really talk about a traffic delivery point per se, but rather of:

- TDDR – Traffic Delivery Destination Region: a region of reachable end-points (which for example may be expressed in terms of a set of IP prefixes). TDDR are used in conjunction with ASQ TT (Traffic Termination, see Section 3.2.6.1) where the destination is a region. All end-points in the region may benefit from the Inter-Carrier good but specific end customer SLAs have to be added to assure that commitments on negotiated performance can be fulfilled during the specific sessions of the individual connectivity service. Thus a region represents only potential destinations, while not all destination points of a region may benefit of the same network performance; for example eyeballs of edge NSPs can have different performances due the used access technologies. This is the case with the ADSL technology for which network performances depends e.g. from the length of the last miles access liaison (between the End-Customer modem and its attached DSLAM equipment). That is why the SLA with Region as destination may include a predefined percentage of the destination points of the region for which the service is quality assured without nominating them individually.
- TDSR – Traffic Delivery Source Region: a region of end-points expressed in terms of a set of IP pre-fixes that is the source of the traffic. TDSR are used in conjunction with ASQ TO (Traffic Origination, see Section 3.2.6.1) where the source is a region. All prefixes in the region may benefit from the inter-carrier good but with the same limitations as TDDR.

In addition to these points and regions that may characterize the ASQ path, we also define specific end points to design the same behaviour between the end customer and the ETICS framework:

- TSEP – Traffic Source End Point: refers to the point where the end customer delivers its traffic to the edge NSP. This can be an end-point that is directly associated with an ASQ path or not. In any case it is up to the edge NSP to decide on the traffic steering and grooming policy. For instance end-user traffic can be steered and groomed onto an interconnect ASQ path.
- TDEP – Traffic Destination End Point: refers to the point where the ASQ traffic is delivered to the end customers. Similar arguments as above apply here as well regarding edge NSP traffic steering and grooming policies.

FIGURE 3 below illustrates the position of the different reference points we have defined around a simple ASQ path. At the ingress side of the ASQ path, end customers send their traffic to their respective TSEP that will be transmitted by the Edge NSP to the TDSP with assured quality. Even if the ASQ Path process aggregates traffic for scalability, it is possible to have a dedicated ASQ path for a particular traffic. In that case, there is only one TSEP. Then, if an ASQ path is split into several segments (from the buyer point of view), the segment boundary is materialized by a Transit TDP. At the egress side, traffic is output from the TDDP to one or many TDEP or to a TDDR region depending on the nature of the contract. Note that an ASQ path at the simplest could be reduced to only one point for Interconnect products. In that case either TDSP is logically co-localised with TDDR, or TDSR is logically co-localised with TDDR, or TDSR is logically co-localised with TDDP. It designs the simple scenario where two Edge NSPs have negotiated a simple ASQ path around a single point of peering.

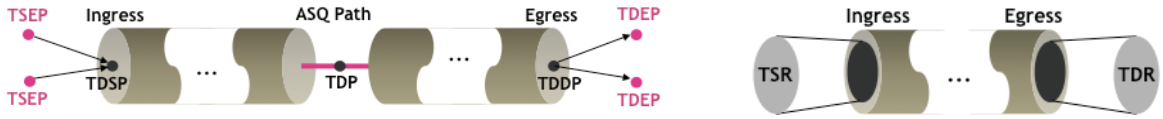


FIGURE 3: ASQ PATH AND LOCALIZATION OF SINGLE OR REGION TDP

FIGURE 4 shows where Transit TDP are localized when ASQ path are stacked to form a hierarchy. In this scenario, the Transit TDP of the lower level ASQ is co-localized with the TDSP of the higher level ASQ, as the higher level ASQ could aggregate several lower level ASQ's for the same destination. At the egress point of the higher level ASQ, lower level ASQ are output like if no hierarchy has been used and again, the TDDP of higher level ASQ is co-localized with the Transit TDP of the lower level ASQ.

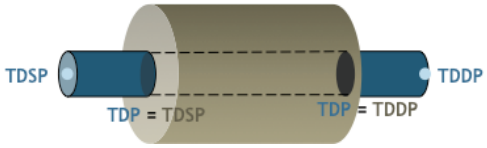


FIGURE 4: TDP AND ASQ HIERARCHY

Segment measurements could take place between two TDPs to verify the conformance of the carriers in terms of SLA contract and QoS. A carrier undertakes to respect and guarantee QoS only between two TDPs, while neither observation nor inspection could be performed inside the carrier network between two TDPs to allow protecting the network operator confidentiality. Carriers must only disclose network information at the TDPs, which represent the public point of information disclosure by the carrier inside the ETICS system. End-to-end measurements could take place between a TDSP and a TDDP regarding an SLA contract

between two NSPs (N-SLA). They could also take place between a TSEP and a TDEP for an SLA contract involving a final ETICS community end-customer (EC-SLA).

### 3.2.5. Traffic Identification for ETICS ASQ Traffic Services at The PoI

Once the Point of Interconnect, Interconnect Interface and Traffic Delivery Point have been defined and selected through the SLA and SLS, i.e. where NSPs exchange traffic, NSPs must agree on how they exchange the traffic, and in particular on how NSPs can identify the traffic at the TDPs in order to properly handle it in their respective part of the ASQ. Again, this kind of information must be precisely defined in the SLS in order for the different NSPs in the ASQ chain to properly configure their respective equipment.

Traffic identification is rather complex and largely dependent on the underlying technology used to guarantee the QoS. In addition, as we are managing inter-carrier and inter-domain traffic exchange, in order to remain scalable, the traffic volume and the number of individual sessions forbid using too sophisticated mechanisms. Therefore, basing traffic identification upon the standard 5-tuplets (source and destination IP addresses, source and destination ports, protocol number) is not possible.

For connection-oriented mechanisms, the solution is rather simple: the MPLS label will identify the traffic that is output at the PoI from one NSP (in conjunction with the ICI of course). If some merging or nesting is performed by the next NSP, label stacking must be used to preserve the original label and to let the different NSPs handle the traffic at the following PoI. Of course, the complex work is deferred to the Edge NSP that must configure the Forwarding Equivalent Class (FEC) on its Label Edge Router (LER) in order to apply the first label to the correct IP packets. The FEC is the standard mechanism defined in MPLS to identify the IP packets that must be labelled. A FEC could be composed by any fields of the IP header and, if needed, part of the protocol (TCP, UDP, ICMP) header.

For connection-less mechanisms, the situation is rather complex compared to the previous one. Again, it largely depends on the underlying technology.

- In IPv6, like for MPLS, the flow label could be used, but without the possibility to stack the flow label.
- For LISP, VPN and GRE, it depends on whether the “tunnel” is terminated at the PoI or not. If it is terminated only in the destination NSP, there is no need to identify the traffic at the PoI. If not, we only could compose with the IP packet header (eventually in combination with the TCP or UDP header), but again keeping the performance and scalability issues in mind.

In case of heterogeneous technologies between NSPs (e.g. MPLS-TE, then VPN, then GRE), traffic identification could again only occur at the IP packet header level.

Therefore, in order to remain scalable, if MPLS or IPv6 label cannot be used, only the *DiffServ* Code Point (DSCP) in conjunction with the destination IP address could be used to identify the traffic at the PoI. Indeed, ASQs are built for a given destination (point or region) and a given QoS. Hence, in order to determine which IP packets belong to which ASQ, the IP destination address and DSCP are sufficient to make that distinction, while remaining scalable as only one Access Class List (ACL) per ASQ needs to be configured.

### 3.2.6. ETICS NETWORK SERVICES ILLUSTRATED AND ELABORATED

Given the ETICS network service types as listed above, this section provides further elaboration and illustration of some of the ETICS network services. In particular, we elaborate on additional constraints that typically must be associated with region based services in order to establish realistic and usable overall solutions. In addition, various aspects of services are identified according to what kind of buyer or requester perspectives are assumed.

#### 3.2.6.1. NSP-to-NSP Network Services

To start with, single-NSP services as illustrated below are considered. The illustrations are assuming TT services. FIGURE 5, upper part (denoted "(a) PoI to Region") shows a service where the hosts are located in the domain of the supplier NSP, which then plays the role of an edge NSP. This is the simplest region-based TT service and can correspond with E1 or E2 reference points. Considering the E1 option, which is considered as applicable in a bootstrapping context, the buyer (customer) NSP is also an edge NSP.

The lower part of FIGURE 5, denoted "(b) PoI to PoEI", shows a point-to-point case where the ASQ traffic is terminated in association with the PoEI. As expressed above, the exact termination of the service associated with the PoEI can vary and hence a more detailed description of the PoEI oriented services are needed in order to achieve the needed precision of the inter-NSP interactions.

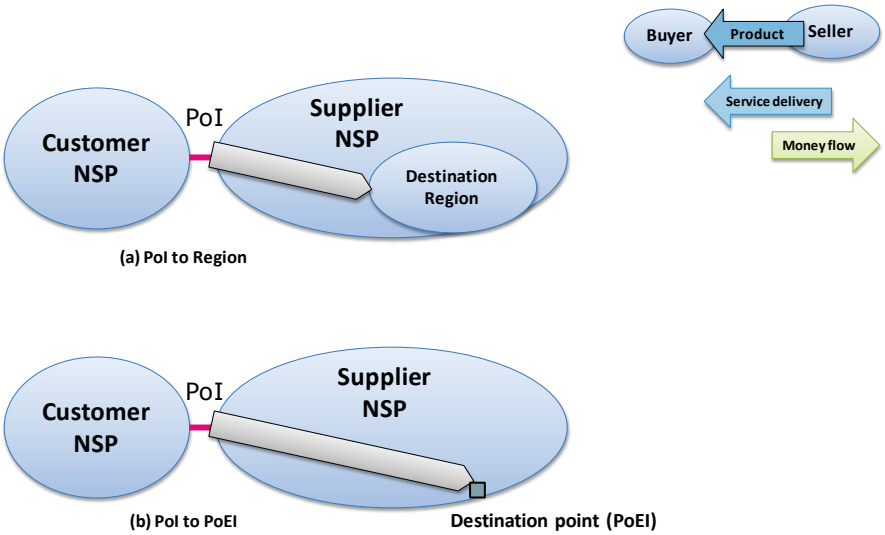


FIGURE 5: SINGLE-NSP SERVICE FROM POI (TT)

Note that in the E2 case, the buyer or the requester may be an entity different from the upstream NSP. This case is not explicitly illustrated as the figure shows that the upstream NSP is the buyer (customer) NSP. This is applicable for both cases ((a) and (b)).

The following two TO based cases are (a) Region to PoI and (b) PoEI to PoI (see FIGURE 6). As described in Section 3.2.1.1, the region based TO service may need to be associated with a corresponding TT service in order to be a viable service. Moreover, additional constraints may then apply on how the traffic is further delivered downstream.

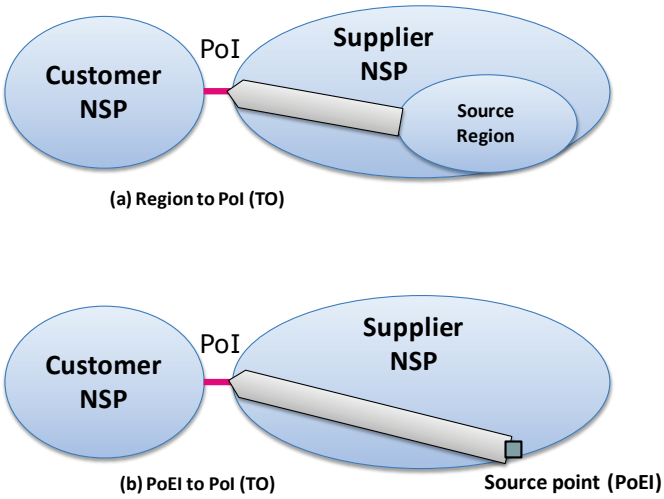


FIGURE 6: SINGLE-NSP SERVICE AT POI (TO)

The next illustrations address the same kind of services as above but this time, the services are multi-NSP services (see definition above). Note that the service as perceived by the customer does not include any information on-, and is independent of-, how the underlying composition is structured or achieved. FIGURE 7 indicates that the service is dependent on a similar service that the supplier is buying from his downstream supplier NSP in the chain. An interesting observation is that while for the case (b), which is a point-to-point path, the downstream service has a 1:1 relationship with this service in focus, while for case (a), the downstream service has an indirect relationship with the service in focus. This introduces the concept of "indirect composition" that will be elaborated in more detail later on.

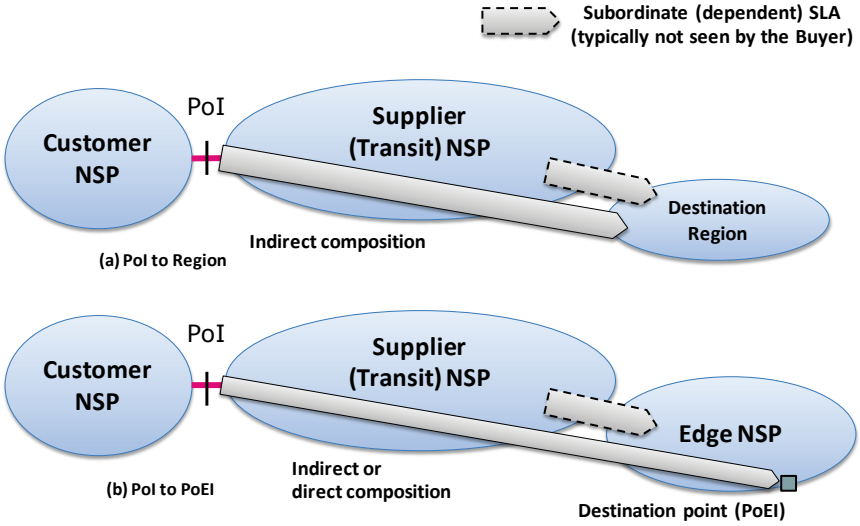


FIGURE 7: MULTI-NSP SERVICE FROM POI (TT)

While FIGURE 7 is showing the buyer (customer) NSP as the upstream neighbour, this does not necessarily need to be the case. ETICS is also supporting cases, where the buyer NSP, or a facilitator requester, is associated with another actor. Again, these options are not explicitly illustrated here. However, the type of network service remains the same at this overall level, whereas there can be service parameters and information elements that can differ from one case to another, depending on who is actually requesting the service.

Taking a closer look at region-oriented services and how they can be composed, FIGURE 8 illustrates an NSP having multiple downstream supplier NSPs as well as multiple upstream customer NSPs. In a similar way, the products and services can go in the opposite direction as well. For simplicity this is not shown, and it is also assumed that one agreement can involve multiple products which again can consist of multiple network services. The interesting observation is, that for region-based services, the general case is that there is an N:M mapping between a downstream service, which is bought, and an upstream service, which is sold. Hence, there is an indirect composition of such services.

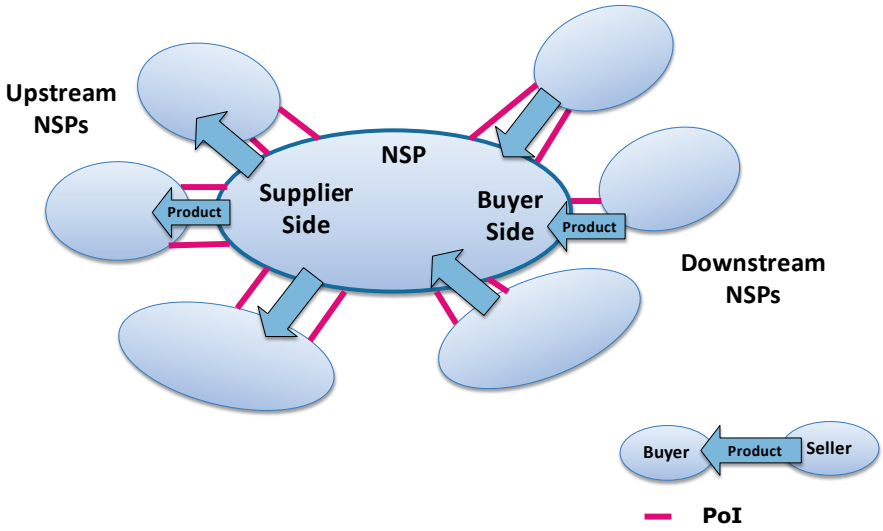


FIGURE 8: CASCADING AGREEMENTS (INDIRECT COMPOSITION) WITH REGION-BASED ASQ TRAFFIC SERVICES

In an early phase of an ETICS solution deployment, it is expected, that this mapping will be done manually and that the ETICS system will manage and control the services on a *"one PoI at a time"* basis. However, as the experience with such services increases, it is expected that the degree of automation can also increase.

One key driver or motivation for region-based services is to enable and support ASQ connectivity for session-based services. In this respect, the region-based ASQ traffic service will need additional constraints and supporting capabilities in order to work as intended. While a region will be defined by a set of IP prefixes, it is not intended that all hosts in the range are active at the same time. Hence, there can be a number of constraints associated with the ASQ traffic service that influence the microflow sessions, including:

- Maximum total traffic (e.g. a scheduled profile per day and hour of the week)
- Maximum number of micro-flows
- Maximum traffic per micro-flow
- Maximum increase in total traffic per 15 min window
- The NSP supplier must always be notified (for information only) by the upstream NSP, which micro-flows have been admitted onto the ASQ traffic service. Hence, it will be possible for the supplier NSP to perform random audit as it finds it is needed in order to check, if the micro-flows have been accounted for at the application plane.
- For a TO service: By random audit it can be verified whether a corresponding micro-flow (opposite direction) exists, and the maximum traffic rate of the micro-flow can be checked.

The random audit approach can be costly, if it has to be used extensively. If low trust level towards a neighbour NSP must be assumed, then the supplier NSP may need to introduce admission control also at the ASBR, if dealing with a session of sufficiently high value; in such a case, the ETICS system can be an enabler for handling the session services. There is no exact general answer, where to draw a line on when to use ETICS for session services and when not.

In addition, the ETICS solution can be used to convey or enable operational support information or management services to manage the session handling features (again, this topic is for further study – cf. the E1' … E3' reference points as described in the Service Enhancement Functional Area (SEFA) part in section 4.1.5 below). One simple example is the information on where to find the applicable SEFA / session handling interface.

Moreover, the region-based service can be used also for non-session host-to-host services. The assumption and constraint can be that the active hosts within the defined range to be operational have been provisioned accordingly at the service edge node of the edge NSPs and that the total traffic volume generated or received by these hosts can be assumed to be within some limits, such that admission control per micro-flow is not needed. Note that some topics or capabilities, such as support for merging or splitting regions, as well as dealing with sub-regions, e.g. for dynamic pricing purposes, are topics for further study.

Next, the single-NSP service from PoI-to-PoI is considered and illustrated in FIGURE 9. In general, the buyer or requester can be a different actor or element than that of the (traffic-wise) upstream neighbour NSP. In this case, the buyer NSP may pay for traffic in both directions.

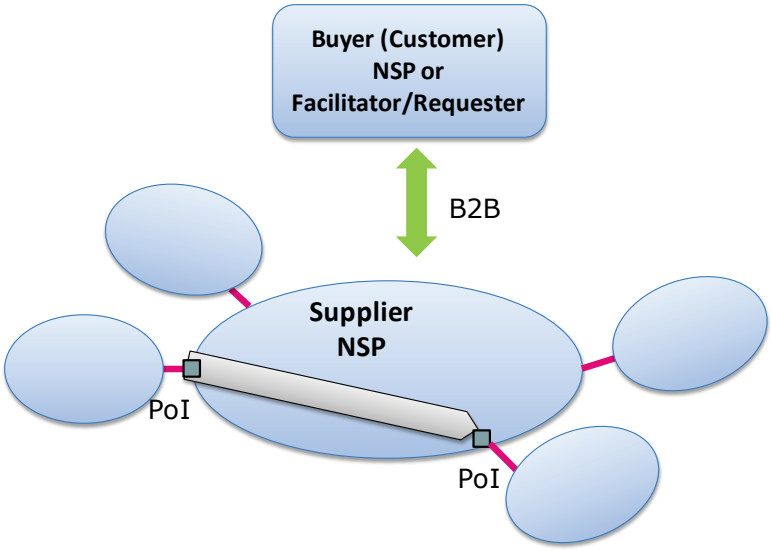


FIGURE 9: SINGE-NSP SERVICE FROM POI-TO-POI (TT)

Similar cases and arguments apply for the multi-NSP service as illustrated in FIGURE 10.

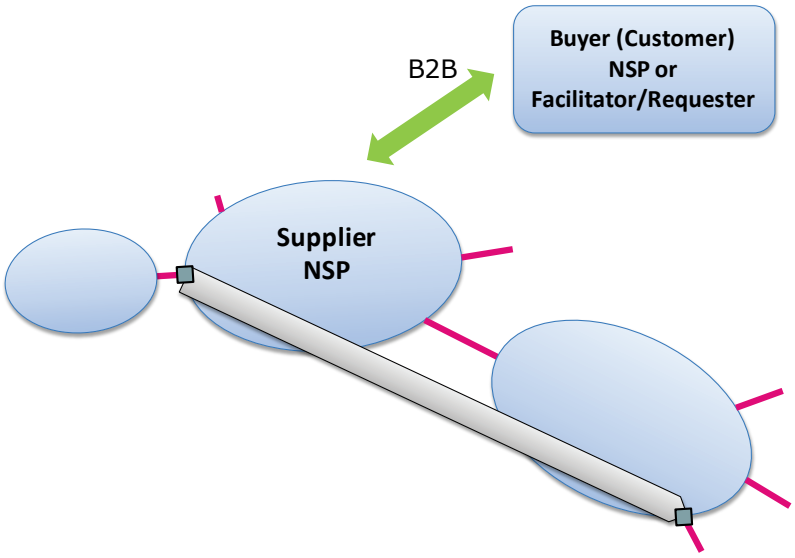


FIGURE 10: MULTI-NSP SERVICE FROM POI-TO-POI (TT)

Furthermore we note that the ETICS network services can be not only composed into a composite ETICS network service, but can also be used in combinations, where the combined ETICS services can be controlled by some relationship constraints. The following case in FIGURE 11 illustrates this where the so-called ASQ TT Adjacency service offered by the transit NSP is enabling the Customer NSP to buy region-based ASQ TT from a virtual remote PoI with supplier NSP D.

In this context there can be a need for the NSP T and the NSP D to interact (using the ETICS B2B interface) in order to decide exactly which ICI and TDP(s) is (are) involved in each specific service. This setting enables NSP B to buy directly from D without having to rely on NSP T delivering the traffic all the way to the destination hosts.

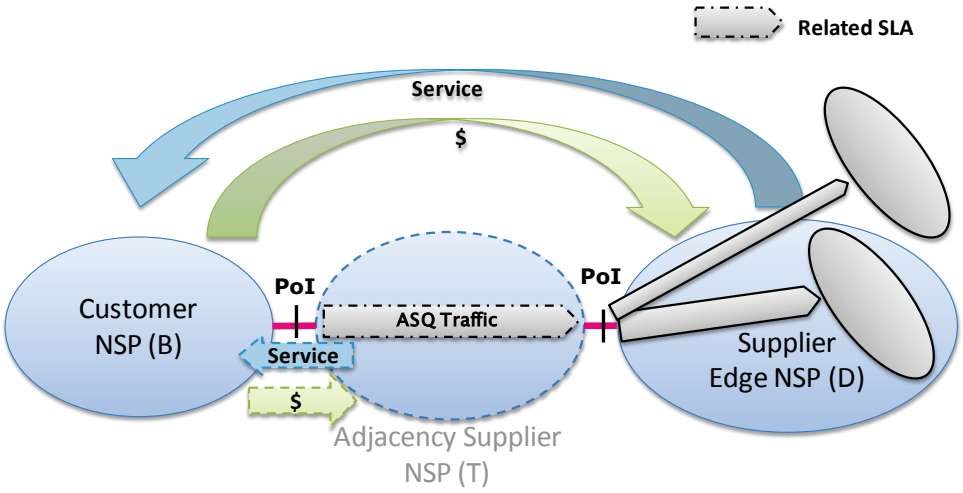


FIGURE 11: EXAMPLE OF ETICS SERVICES IN COMBINATION: ASQ TT BY MEANS OF ADJACENCY OFFERING

### 3.2.6.2. NSP-to-Enterprise Network Services

A PoEI-to-PoEI ASQ traffic path can actually be realized in many ways, and such an ETICS network service can be terminated in many ways, i.e. on customer side of PE, on CE, or on NE of the Enterprise Network.

FIGURE 12 illustrates an end-to-end PoEI-to-PoEI bidirectional ASQ traffic path offered to an enterprise customer. The enterprise customer only sees the interconnect end-points (PoEIs or end-points resolved into more specific end-points, cf. TDEPs/TSEPs above) associated with his two enterprise networks E1a and E1b, respectively. Such an ETICS network service can be realized in many ways. However, these details are hidden from the enterprise customer. Here the PoEI-to-PoEI ASQ path is shown as a concatenation of three segments across NSP A, B and C, where the NSP A is facing the enterprise customer with an ETICS compliant B2B service management interface. NSP A has an internal service management interface as well as ETICS B2B interfaces with NSP B and C that enable the composition of such a service.

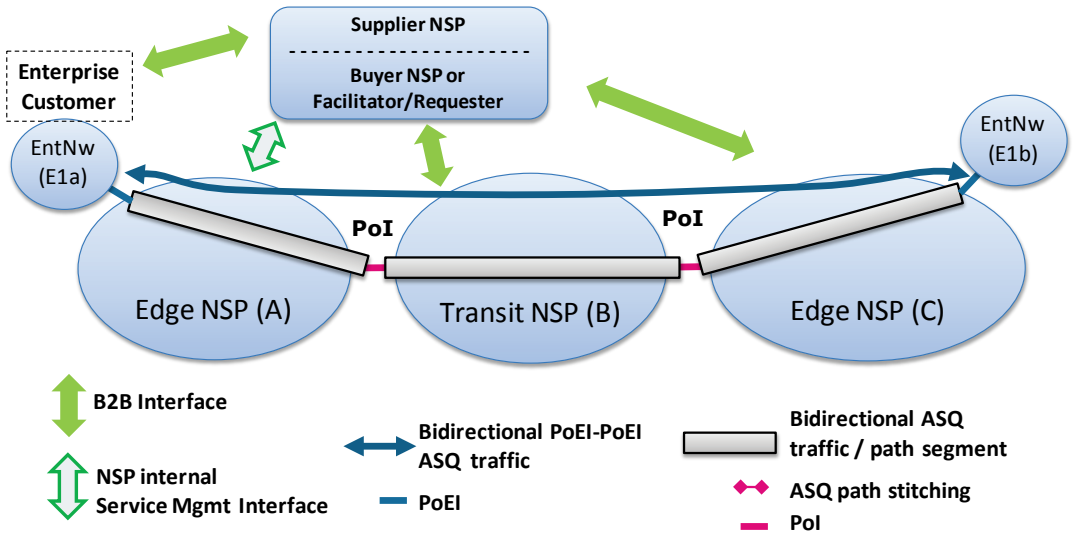


FIGURE 12: BIDIRECTIONAL ASQ PATH FROM POEI-TO-POEI

In addition to the above way of realizing the end-to-end service, where each ASBR on the path must have policies configured for the path, it is also conceivable, that the end-to-end service can be based on region-based ASQ TT services in combination with ETICS facilitated configuration of the PoEIs, in order to avoid enterprise customer state in the ASBRs. The preferred way of realization will depend on the service requirements.

Furthermore, inter-NSP interworking for customer VPN routing information dissemination and operations support can also be supported by an ETICS based approach. However, these topics are for further study.

In the next example, an ETICS service is offered to an enterprise customer – here an Over-the-Top (OTT) information service provider – which is offering OTT content-distribution-network (CDN) service to his content provider customers (not shown). The OTT is buying ASQ TT from the edge NSP1 enabling delivery of content session services to hosts within the destination region as illustrated in FIGURE 13. However, this ASQ TT service must be supported by a service enhancement and session handling functionality as indicated in order to assure resource and admission control or any other supplementary service for the given end-user requesting the content delivery service. The way the end-user requests the content service is outside the scope of ETICS.

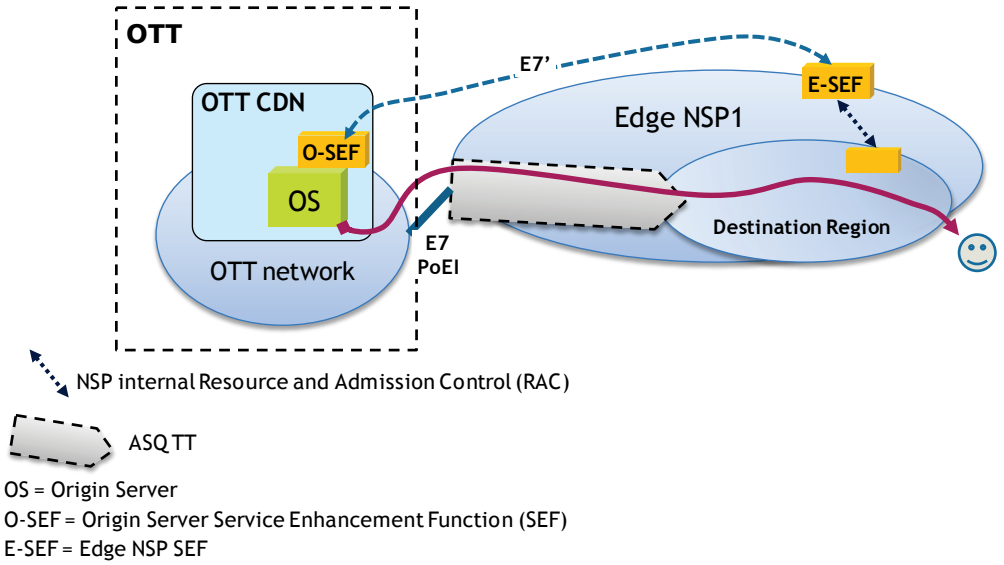


FIGURE 13: CONTENT DELIVERY BY MEANS OF ETICS SERVICES

In a similar way, ETICS services can be used to enable a Telco SP operating a Telco CDN solution to offer content delivery services to Business Customers. The usage of inter-NSP ETICS services for the support of CDN interconnect connectivity is a topic for further study.

The above illustrations provide a few examples of ETICS services offered over the E6 and E7 reference points (potentially supported by E6' and E7', see the SEFA topics in Section 4.1.5 below). The various settings to be taken into account for the support of business cloud connectivity services have not been specifically addressed here, but represent great opportunities for the ETICS system framework.

### 3.2.7. Business and Technical Parameters of ETICS Products and Services

An NSP product can be either offered or requested. Each NSP product has parameters that specify it. We now present these different parameters and separate them into business parameters and technical parameters.

#### 3.2.7.1. Technical Parameters

##### *3.2.7.1.1. QoS Parameters*

ASQ product offers are differentiated by the level of network performance they assure. We define the well known set of QoS parameters as follow: Delay (**D**), Jitter (**J**) and Loss[18] (**L**). These parameters characterize the QoS performance an NSP provides between the demarcation points at the entry and exit of its domain. The basic parameter of an NSP product is represented by the Bandwidth (**B**), which corresponds to the amount of traffic which can be booked for a service request instance. We can imagine that NSPs define minimum, maximum and intermediate step values to fix acceptable bandwidth to request or order.

In addition to these parameters, a crucial criterion for the quality of a network service product is its Availability (**A**). It could be used by an NSP to select a dedicated technology/mechanism (e.g. add path protection, fast re-route …) and a device configuration (e.g. which queuing mechanism will be used).

[18] Can be expressed as packet loss or bit loss.

QoS parameters defined in the SLSs may have dependencies between them. In particular, the availability may be constrained by the fulfilment of other parameters {D, J, L}.

In the following section of the deliverable, we will use the notation Q = {D, J, L, B, A} to denote the QoS objectives of an SLA contract.

*3.2.7.1.2. Time Parameters*

In addition to technical parameters that describe mostly the QoS, the SLS could also embed some parameters to precisely define from which moment and for how much time the SLS is requested. First of all, we define the Time Duration (**TD**) as the duration of the service that can be ordered. To ensure minimum network stability and also indirectly reduce scalability issues, acceptable values for time duration can be specified with a minimum value and a given time interval for incremental steps. In addition, a maximum duration may optionally be specified. We also define a Delivery Delay (**DD**) for the offer as necessary to make the service available. In most of the cases, DD expresses the maximum time to perform a network element configuration in an NSP domain for provisioning the service. But **DD** can also represent the delay to obtain this service when NSPs are out of stock of some resources (i.e. no more bandwidth is immediately available). **DD** could also be the default behaviour when new installations are necessary, such as adding an optical fibre link, setting up a new optical wavelength in an optical switch, etc. It could finally be used to plan reservations in advance. In the requested SLS, the DD parameter is interpreted as the maximum amount of time accepted by the customer for the service activation. Finally we introduce also the parameter start time (**ST**) to specify, in the requested SLS, when the connectivity is supposed to be established. The ST parameter differs from DD's usage in a request. ST is used to schedule a service in a date in the future; the service must not begin before. DD only stipulates that the service must not start after a date. If DD and ST are both used in a request then the date defined with using the DD value must not be before the date defined using the ST value. With TD and ST, an End-Customer or an NSP could plan its SLA in advance (scheduled service).

An additional benefit can be derived from the use of TD: it can be used as a means to protect against dummy reservations. Indeed, if a failure appears in the system, there is a great chance, especially if service termination synchronisation failed or is not supported, that a given SLA contract will be lost, and thus never properly terminated. To avoid this, each SLA must be associated with a timeout, i.e. a validity time after which the SLA is automatically removed. Even if the customers (both EC and NSP) are not specifying the parameter TD, a default value will be assigned to the contract (based on the validity of the offer – see below).

#### 3.2.7.2. Business Parameters

Quantitative business parameters are the price (**P**) and the cost (**C**). If the price P is disclosed in the offer, the cost C remains an internal value for the NSP. In fact, NSP could obtain from the control plane (or other means) a cost C for a given Q = {D, J, L, B, A}, i.e. C = f(Q) where f is a cost function. From this cost C, the NSP could apply some internal policy and business rules, i.e. different algorithms, in order to compute an acceptable price P, i.e. P = g(C) where g is a price function. Other parameters could also be taken into account, but at least the difference P–C will give an indication of the revenue an NSP could expect from the service. The cost, and thus the price could be dependent on the amount of bandwidth, the duration of the service, the start time, the availability of network resources, etc. It is reasonable to think that an NSP will

make some discount offers regarding the volume and the duration, e.g. price per Mbit/s per month will be cheaper for 1 Gbit/s for two years rather than 100 Mbit/s for one month. But the price is not only a factor of B*TD. For example, a 40Gbit/s connectivity service for only one hour will be certainly much more expensive than a 1Gbit/s connectivity service during 40 hours. Indeed, the 40Gbit/s service will have a stronger impact on a network domain than multiple flows which can be distribute in time and in space. Therefore, for such big throughputs, NSPs will encourage longer service time durations.

#### 3.2.7.3. Validity of Network Service Offers

The Mescal project [BoLe05] and IETF [RFC5160] have identified a risk to freeze the QoS market when bundles (product offers combining several product offers potentially from different NSPs) are used. In fact, if an NSP A proposes an offer O1 to an NSP B, NSP B could use this offer O1 to build on its own new offer O2, and so on. The main problem arises when the first NSP A wants to change its offer O1, as all other offers using it will be affected. To avoid this problem, we define a Validity Period (**VP**) that (more precisely) defines for how long an offer is valid and can be ordered by a customer. After this period, the offer will be automatically removed or modified by the NSP. In turn, another NSP will not be able to use this offer in a bundle or to publish a new offer aggregating it, after the validity period. Therefore, $VP_{Ox} \leq Min \{VP_{Oi}\}$, where Ox is a new offer built on multiple offers Oi.

To assure the stability of the ETICS system, WP5 studies will have to estimate if a minimum duration for the validity period has to be imposed for all product offers and if synchronization in the offer publication is required in some push scenario. Impacts of such rule on aggregated offers have to be analyzed to determine the viability of the system.

## 3.3. END-TO-END ASQ AND ADMISSION CONTROL FOR END-USER CONNECTIVITY SESSIONS

One goal of this section is to point out how the basic (aggregate level) ETICS ASQ traffic services may be augmented with parameters and capabilities that enable different end-to-end QoS strategies for instance in relation to interconnect traffic steering policies and inter-NSP coordination of admission control.

It is recognized that the need for inter-NSP or NSP-to-Enterprise B2B automation and collaboration in relation to these services capabilities is also an important area. Areas to consider more specifically are provisioning, assurance, statistics/usage tracking, monitoring, and product/service configuration, but most of these topics are not within the scope of this deliverable.

In this section we consider cases where end-user connectivity session service establishment and admission control is needed. We also assume that the ASQ path infrastructure (aggregate level) and business agreements have already been established prior to the time there is a need for end-user connectivity session service establishment. These are represented by the interconnection SLAs between the NSPs. As such, the parts of the ASQ path infrastructure that are controlled by one provider may or may not be visible to other providers.

While the focus of this section is session handling for the purpose of service establishment and admission control, these session handling capabilities are to be considered in the context and part of a wider concept, that is, the so-called Service Enhancement Functional Area (SEFA), presented in Section 4.1.5. Note also

some discussions in Section 3.2.6 related to how region-based ASQ traffic services can be associated with various constraints in order to make them feasible for supporting ASQ session services.

### 3.3.1. General Assumptions

For the general discussion here we assume that admission control is only needed in the edge NSP domains involved in the end-user connectivity sessions and that the source edge NSP has sufficient knowledge about the current availability of the aggregate level interconnect resource towards the other (destination) edge NSP. This is illustrated in FIGURE 14 below, where we assume that somehow the edge NSP has knowledge about his available resource toward a destination region. Again this knowledge is based on the SLA between the actors and the ASQ path infrastructure, which may be either connection oriented or connection less, so in principle it is only the SLAs with the other operators that are relevant. Note that in more advanced cases admission control may also be needed in transit. In such cases this will be reflected in the SLA with the transit NSP.

In addition towards some destinations, it might be that a new end-user connectivity session demand will:

i) trigger a new ASQ path establishment;

ii) trigger an ASQ path modification; or

iii) only trigger set-up of an end-to-end connectivity session without pre-existing ASQ path to establish the session on demand. Such more advanced cases are out of scope of this deliverable.

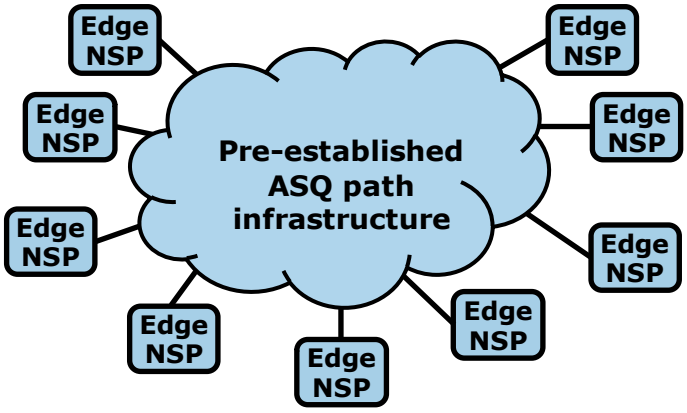


FIGURE 14: EDGE NSPS CONNECTED WITH PRE-ESTABLISHED ASQ PATH INFRASTRUCTURE

The general approach could be that each edge NSP has an SLA towards neighbouring NSPs (either directly with another edge NSP or with a transit NSP) guaranteeing a certain bandwidth with a certain quality towards a certain destination range. The “size” of the destination range will be defined by the downstream NSP (supplier role) and may for instance be a range of destinations within an edge NSP, a range of destinations corresponding to an entire edge NSP, or a set of edge NSP destination ranges. A set of such SLAs would be needed to cover different quality classes and different destination regions, and where relevant also specify per interface if the upstream NSP has more than one interface towards the

downstream NSP. When setting up an end-user end-to-end connectivity session, several signalling sequences may be used. Typically, we can have a case where the upstream edge NSP is responsible for:

- the quality and admission control from the sending party to the PoI with its interconnecting downstream provider;
- choosing the correct quality class and traffic delivery point in accordance with the SLA with its interconnecting downstream provider such that the composite e2e requirements of the session are respected;
- setting the correct packet markings for this quality class (e.g. DSCP value, MPLS label, etc. as appropriate) at the PoI between the two neighbouring interconnecting providers;
- respecting the bandwidth limitations given by the SLA with the interconnecting downstream provider for the given quality class and destination region to which the other end-point belongs;
- coordinating by some means admission control with the other edge NSP. In the simplest form this can be just checking whether the other edge NSP is able to accommodate the connectivity session service for the end-user. In more advanced cases, some negotiation can be needed if there is a need to adjust the level of committed resources in an inter-NSP coordinated fashion.

The downstream NSP provider may be responsible for (in the relation to the upstream provider):

- the quality from the PoI to the end-user.

A transit operator has both, the role of a downstream provider and that of an upstream provider, and different sets of requirements apply in the different roles.

Note that the means and strategy of performing admission control inside an NSP is an internal matter to the NSP and several technical solutions as well as approaches can be envisaged. For instance, admission for the core segment of an NSP can be conditionally performed and only in case there is an event triggered indicating that a higher probability of experiencing congestion on certain core segment, there is a need for performing admission control on that specific core segment.

According to the division between the ***network planes*** accommodated by the NSP role, and the ***application planes*** accommodated by the InfSP role or the Telco application SP role, the session handling can be limited to addressing the connectivity session handling only. For the purpose of session handling, the ETICS architecture foresees a Session Handling Function (SHF) in accordance with the SEFA concepts introduced in Section 4.1.5.

ETICS does not intend to “reinvent the wheel” in this respect, but rather show, how existing solutions can be used and where needed, how existing solutions may need to be complemented or extended in order to achieve a consistent end-to-end QoS approach across NSPs. Hence, the ETICS approach will be compatible with for instance IMS, SIP and Diameter based protocols as needed. Although not further elaborated in this section, different charging approaches should be supported, but only limited to a set of strictly needed capabilities in order to keep costs for charging low. Again, the division between the ***network planes*** and the ***application planes*** should be taken into account.

### 3.3.2. Two Edge Providers with One Logical Interconnect Link

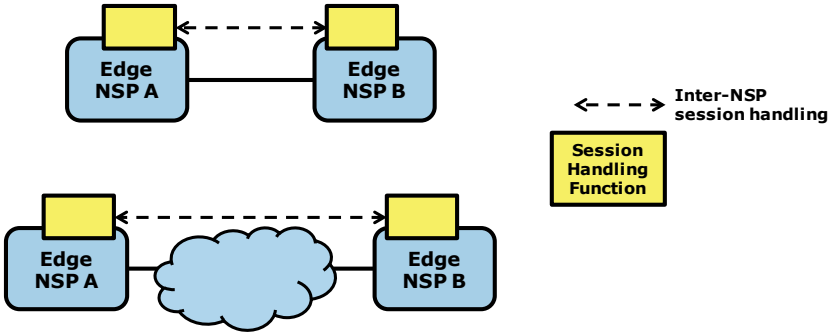


FIGURE 15: JUST ONE LOGICAL LINK BETWEEN EDGE NSPS. NO NEED FOR TRAFFIC STEERING POLICIES

A constellation of two edge NSPs (here NSP A and NSP B) being interconnected via a single logical interconnect link represents the simplest case, which is considered. Some coordination is needed between the two edge providers as provided by the Session Handling Function shown in FIGURE 15. This could be related to session negotiation (e.g. SIP, IMS) and to admission control in both ends (e.g. RACS).

External partner may execute Session handling as shown in FIGURE 16. In such case external partner need to communicate across the vertical Service interface (API) with the NSP, and via its Session Handling Function for admission control and for coordination between the two sides.

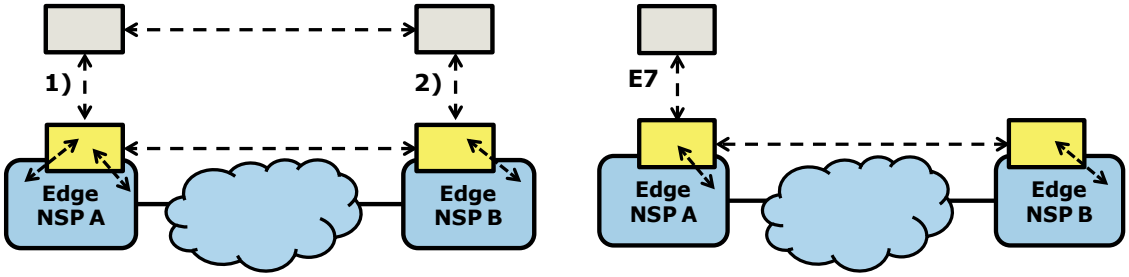


FIGURE 16: JUST ONE LOGICAL LINK BETWEEN EDGE NSPS. SESSION HANDLING BY EXTERNAL PARTNER

### 3.3.3. Traffic Steering Policies

With more than one IC link (ASQ path, logical or physical) between two neighbouring NSPs, the SLA between those may include traffic steering policies, i.e. to be able to satisfy given (QoS) requirements towards certain destination regions for some traffic, the downstream NSP may require that the traffic is routed over one (or more) specific interfaces, and not routed over the other interfaces. Thus, in FIGURE 17 NSP T1 may only guarantee a certain quality level towards NSP B over one of the IC points with NSP A, and NSB B may only guarantee a certain quality level towards some destination ranges over one of its IC points towards NSP T2.

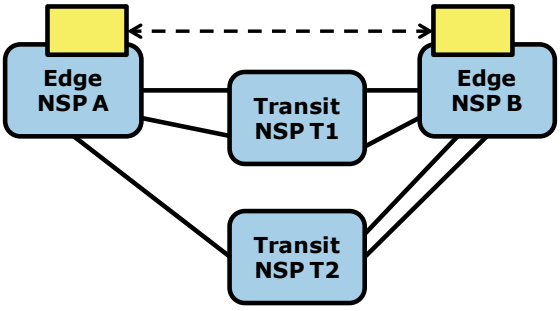


FIGURE 17: WITH MORE THAN ONE LOGICAL LINK BETWEEN ADJACENT NSPS TRAFFIC STEERING POLICIES MAY BE NEEDED

Traffic steering policies are normally semi-static. However, dynamic traffic steering policies may be an option. That is, NSP T1 (in FIGURE 17) may then dynamically notify NSP A of changes to traffic steering policies. These policy changes may be a result of dynamic changes to the traffic steering policies/ agreements with remote downstream NSP. This may result in traffic routed via another transit NSP.

Update of traffic steering policies may be a result of changes in link utilisation. High capacity links will normally not need any Admission Control (AC) to support a given quality level, i.e. to support ASQs. But if link utilisation crosses a given provider-defined threshold, this may activate AC for this link. Also link failure will of course have effect on traffic steering and routing in general, but this is not further discussed here. To discuss the effect of link status changes on traffic steering policies we look into some distinct cases. When we talk about a "link status change", we refer to a case where AC needs to be activated on this link. Only changes that impact the Interconnect, that is, the ability to support a given ASQ connectivity session over a given IC link or ASQ path, will be discussed here.

Many cases and inter-NSP control and signalling (messaging) procedures can be envisaged, and if sophisticated or very demanding scenarios must be supported one may rather consider establishing dedicated edge-to-edge ASQ paths between the edge NSPs which enables a stricter control. The approach envisaged here aims to take advantage of and develop a simple, and in most cases sufficient, ASQ interconnect approach.

While the below cases suggest that the traffic steering policy update notifications can be handled by the Session Handling Function, we recognize that the exact mechanism that should be used for such kinds of SLA/SLS updates and information propagation can also be a topic for or a task of the network service and business plane. The Session Handling Function can then be updated accordingly to obey the new traffic steering policies and behave according to the updated network state information.

Case A: No transit. Any link status changes will be handled by Session Handling Function. Part of this may be choice of IC link for a communication path if multiple links are available between the two operators.

Case B: Communication path with one or many transit hops (e.g. FIGURE 19) and link status change of a link in the NSP A domain that effect at least one of the routes from NSP A towards NSP B domain. Session Handling Function will handle this.

Case C: Single transit hop and link status change of a link in the NSP B domain (FIGURE 18) that affect at least one of the routes from NSP A towards a destination region within NSP B domain. Normally AC can handle this for the involved link and we do only need to involve the Session Handling Function (accept or reject

request). In some cases, the problem can be handled by doing a traffic steering update such that the traffic is routed via another IC link to NSP B. This may have the consequence that the traffic will be routed via another transit NSP (i.e. use NSP T2 instead of NSP T1).

Case D: Single transit hop and link status change of a link in the NSP T1 domain (FIGURE 18) which affects at least one of the routes from NSP A towards a destination region within the NSP B domain. In this case, there are two possibilities: If some other IC links can be used for this communication, a traffic steering update is necessary and notification about this needs to be sent to NSP A and/or NSP B as appropriate. The other possibility is that the transit NSP indicates that AC is needed for this transit and a notification from this provider that it needs to take part in the session handling is sent to the appropriate NSP(s). This is discussed in the next paragraph.

### 3.3.4. TRANSIT NSP TAKES A ROLE IN INTER-NSP SESSION HANDLING

Inspired by the existing hub models in the industry, it is expected that new or evolved hub models can be developed. This will enable a transit NSP to offer hub (proxy) functions such that the customers (Edge NSPs or InfSPs) of the hub NSP can indirectly (via the hub/transit NSP) establish relationships with a set of other Edge NSPs and thus can take advantage of already existing business relationships. Such hub or transit actors may provide functions and services that go beyond the ETICS scope, for instance in relation to the application layer.

In general, a transit NSP that takes a role as Session Handling Hub is (typically) not involved in admission control. The transit NSP may take a role in the routing of the inter-NSP session handling/admission control messages, and might potentially also take part in the charging process, in which case it could enable cascaded charging. However, charging is not further addressed here.

In certain cases, the transit NSP may however require that it must also take part in the admission control. The decision of whether or not the transit NSP takes part in admission control may or may not be exposed to the other NSPs and different signalling strategies may be envisaged.

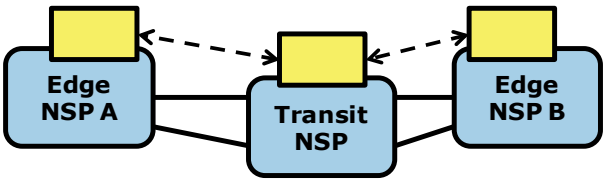


FIGURE 18: TRANSIT NSP TAKES A ROLE IN THE INTER-NSP SESSION HANDLING

In FIGURE 18, the simple case with only one transit NSP involved is shown. In this example, due to high load on (a) certain link(s), the transit operator has decided that AC is needed to set up an ASQ session between edge NSP A and edge NSP B via the transit NSP domain. In the figure, it is indicated that NSP A first checks with the transit NSP. If ok, the request is forwarded to NSP B. Another possibility could be that NSP A first communicates with NSP B before involving the transit NSP in the session handling.

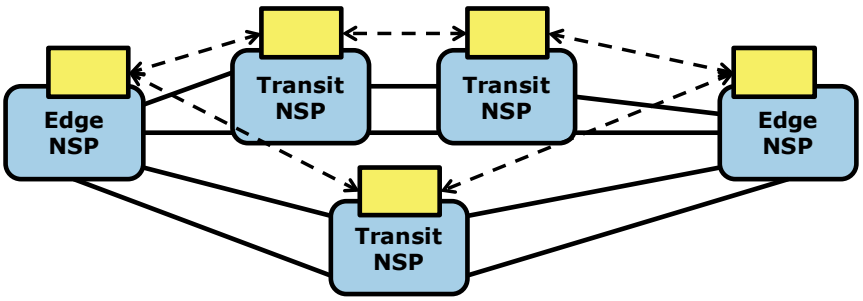


FIGURE 19: TRANSIT NSP TAKES A ROLE IN THE INTER-NSP SESSION HANDLING. MORE COMPLEX SCENARIO

More complex scenarios involving a chain of transit NSP that take part in the end-to-end admission control can also be envisaged, as indicated in FIGURE 19 above. These scenarios are for further study.

## 3.4. ETICS POLICY RULES

An "ETICS NSP" is an NSP that can offer ETICS network services according to ETICS specifications. Today, NSPs are just interconnecting their respective Data Plane and Routing Plane (through BGP) at the peering point. First iterations of ETICS architecture (D4.2) and requirements (D2.2) have explored ways to interconnect their respective Business/Service Plane and Control Plane in order to offer more advanced network services. These kinds of interconnections may not only be governed by the current standard protocols. The set of ETICS NSPs as well as NSPs acting as buyers of ETICS network services constitute what is referred to as the "ETICS community". Hence, the notion of ETICS community is flexible in terms of how the NSPs interact and do business. Once the EU FP7 Project ETICS terminates, it is foreseen that ETICS specifications can be handed over to a new or existing forum or standardisation body.
In addition, it is foreseen that various types of associations, federations or alliances with additional governing policy rules for its member NSPs can be formed. Thus, to enhance the "ETICS Community" description, some initial rules are suggested and discussed. These rules are targeted to protect NSPs from undesired behaviour and to avoid any particular NSP becoming predominant in the community. These rules, tentatively named ETICS Policy Rules in the following, are divided in two categories: Technical and Business Policy Rules.

This section presents work in progress on such policy rules and various example types of associations or alliances that could fit our telecom market requirements. Such policy rules have already been studied in past projects. The various types of associations, federations or alliances can range from open ones with few rules and constraints on membership (e.g. open association) to acceptance-based ones with a strict set of rules where membership is by invitation only (e.g. closed alliance).

### 3.4.1. STATE OF THE ART

This sub-section provides an overview about past collaborative projects that have conducted research on telecom markets and which propose definitions for alliances and federations of network operators.

#### 3.4.1.1. ACTRICE Alliance

The French national project (RNRT), named ACTRICE, is one of the main activities on inter-carrier services preceding the ETICS project. In [ACTRICE-D1.3],an alliance is defined as a cooperation agreement in which two or more institutions work together in order to mutually share resources / inputs while preserving their

own identities each. An alliance implies a degree of strategic and operational coordination. It may take the form of a joint production or marketing, exchange of technology, know-how, etc. Alliances can serve several purposes, such as reducing risk and uncertainty, sharing R&D cost, achieving economies of scale, opening up access to technologies and markets, etc.

An alliance is an intermediate or hybrid form of the market (short-term contract between two partners) and hierarchy (integration of partners within the same company). In an alliance, the partners remain independent, but bound by a contractual arrangement for the long term.

ACTRICE supports alliances as a solution to enhance cooperation and trust between operators and service providers. Among the objectives of an alliance between Internet Service Providers (ISPs), ACTRICE has mentioned the support of new services with QoS guarantees. The alliance can also manage the uncertainty and incompleteness of contracts and protect itself against the risks of opportunism. Finally, the alliance facilitates standardization and interoperability. The latter is a major motivation for the formation of alliances between operators to set up an infrastructure capable of supporting and developing QoS services such as defined in ETICS through the service plane.

Based on a comparative analysis of two market sectors already using alliances, the air transport sector and industry credit cards, ACTRICE provides the answers to the first question concerning the choice of partners: What are the companies like that will join the alliance? And also, once the first question has been resolved, which coordination mechanisms must operators agree on? In other words, what types of contracts do traders use in order to formalize their alliance and what is the form of this alliance (centralized or decentralized)?

In summary, three conclusions have been drawn. First, ACTRICE has highlighted the importance of alliances as a way of organizing economic activity and value creation. Through an alliance, fixed costs can be reduced and the quality of service improved. The payment card industry may serve as a good example in favour of the claim that centralized alliances offer an attractive mode of coordination for an infrastructure with QoS.

#### 3.4.1.2. DAIDALOS Federation

FP6 DAIDALOS-II project has studied and proposed a federation model for operators who wish to address the inter-carrier market. They have identified three kinds of federation:

- Data federation, where only provider-specific data are exchanged which are not directly related to the end-users,
- Identity federation, where the identities of customers are exchanged between carriers,
- Function federation, where functionalities provided by the different carriers are exposed to others.

Concerning hierarchy, two different types of federations have been envisioned by DAIDALOS:

- Horizontal federation is established when two or more providers wish to exchange traffic for a particular service on the same communication layer (e.g. mobile roaming),
- Vertical federation is established when two or more providers wish to collaborate across different communication layers in order to provide a value-added service.

In DAIDALOS terms, ETICS federations could primarily correspond to the types *Data* and *Horizontal*. Vertical federations could be of relevance for the NSP and the Transport Provider when making use of the E4 interface.

### 3.4.2. FLEXIBLE GOVERNANCE

A particular "Policy-governed ETICS community" is defined as "a group of NSP that have common interests in the telecom market". In terms of products, this corresponds to common interests in business objectives and acceptable business rules in the market. In particular, it is a market place that intends to facilitate the publication, composition, negotiation, monitoring and exchange of Service Level Agreements (SLAs). The goods exchanged within the "ETICS community" are assumed to be Assured Service Quality (ASQ) Inter-Carrier (IC) connectivity and traffic.

Depending on the level of constraints imposed by the ETICS Policy Rules, an ETICS community could be:

- An **Open Association** where no policy rules are on place or at least very flexible and do not bind its members to strict constraints. In this case, the ETICS Community is "*a market place where an NSP can choose to participate by selling or buying ASQ goods according to the ETICS specifications*",
- A **Federation** when Technical Policy Rules require strict behaviour but where Business Policy Rules remain flexible. Inside the Federation, a certain level of trust could emerge between NSPs and some common SLAs could be negotiated,
- An **Alliance** when both Technical and Business Policy Rules require strict behaviour. The Alliance implies a high level of trust between NSPs, e.g. to allow shared SLAs and penalties.

#### 3.4.2.1. Joining an ETICS Community as Member

Each Network Service Provider that explicitly asks for access to an ETICS community could potentially become a member. Before the NSP joins the community, it must accept all Policy Rules (which apply to its membership level) that govern the community. Once the NSP has joined the community and becomes a member, it can start to operate, i.e. sell and buy ASQ IC goods to/from the other NSP within the community.

#### 3.4.2.2. Withdrawal from an ETICS Community

Each Network Service Provider that is member of an ETICS Community can withdraw at any given moment provided that it has released all ASQ IC goods (both the ones it uses and the ones it provides). In case an NSP does not respect one or more Policy Rules that govern the community, it will be evicted and all ASQ IC goods that are linked to this NSP will be removed.

### 3.4.3. ETICS POLICY RULES DEFINITION

This section aims to define a first set of tentative policy rules that could govern the different ETICS Community types (i.e. Open Association, Federation, Alliance) and that NSP must follow and respect if they want to be part of the community. Policy rules are subdivided into two categories, i.e. technical & economic ones.

#### 3.4.3.1. Guidelines and Common Definition of the Policy Rules

**Announcement:** This represents an SLA offer or Network Capabilities advertised to other NSPs. An announcement is composed of a certain number of information elements as defined in [ETICS-D4.2]. At least, it must carry one or more network prefixes and associated QoS objectives that the NSP will guarantee when receiving traffic for these prefixes from another NSP. By using the CIDR notation, an IPv4 network prefix could designate a customer (/32) or a region (/x where 0 < x < 31). An announcement always remains the property of its creator.

**Public Announcement:** An announcement destined to all NSP within the alliance.

**Private Announcement:** An announcement destined only to a given NSP (the receiver in a peer relationship).

**Re-bundled Announcement:** An announcement that could be re-bundled by another NSP.

Note that Re-bundled Announcements can be public or private.

#### 3.4.3.2. Technical Policy Rules (TPR)

The following technical policy rules are provided following the SLA life cycle defined in deliverable D4.2. In each table, the third column will define to which level of association the Policy Rule may apply (O = Open association, F = Federation and A = Alliance).

**SLA Offers and Network Capabilities Discovery**

The following technical policy rules govern the first steps of the SLA life cycle and in particular the publishing phase. The main objective is to provide the same visibility to all NSPs within the alliance about what is publically available both in terms of SLA offers and Network Capabilities and to avoid that a particular NSP retains any announcements.

| TPR# | Description | Type |
|---|---|---|
| **TPR1** | All NSPs must send **announcements** about their networks within the ETICS community. | O, F, A |
| **TPR2** | All NSPs must inform the other NSPs within the ETICS community (through advertisements) about their Point of Interconnect (PoI). | O, F, A |
| **TPR3** | **Public Announcements** must be forwarded without any modification by the receiver NSP to all its neighbours. In fact, public announcements must be flooded within the ETICS community and all NSP must participate in the flooding process. | F, A |
| **TPR4** | **Private Announcements** must be kept by the receiver NSP and they must never be disclosed to other NSPs. | O, F, A |
| **TPR5** | Only **Re-Bundled Announcement(s)** may be bundled by an NSP with its own announcement(s). | O, F, A |
| **TPR6** | All **announcements** must be sent only to the Facilitator in the fully centralized scenario. | F, A |

**Negotiation and Composition of ASQ**

During the negotiation and composition phase, NSPs must follow the rules specified below.

| TPR# | Description | Type |
|---|---|---|
| **TPR7** | NSPs are free to choose the initial scenario (Push, Pull, Per-NSP, Distributed) to perform the service composition. | O, F, A |
| **TPR8** | The Facilitator in the fully centralized model must answer all NSP requests. | F, A |
| **TPR9** | When PoIs which are to be used are pre-determined, the NSPs must follow the given PoI when it performs the service enforcement. | F,A |

The topics of restrictions and flexibility of which negotiating and composition method to follow will be considered as various types of ETICS communities are further analysed.

**Monitoring and Penalties**

During the SLA lifecycle, once the ASQ is in place and used, monitoring of the SLA is done in order to verify if the QoS commitment is effectively met. NSP(s) violating the committed QoS could be detected in order to both recover the agreed QoS levels and to apply penalties. Instead of – or in addition to – penalties, a "reputation system" could be used.

| TPR# | Description | Type |
|---|---|---|
| **TPR11** | An NSP must deploy a monitoring system compatible to the specification of the ETICS. | F, A |
| **TPR12** | An NSP must announce its (technical) monitoring contact point (M-Proxy). | F, A |
| **TPR13** | If monitoring is done by the NSP offering the product (ONSP), all other NSPs involved in this SLA must allow the ONSP to retrieve their monitoring data to the extent defined by the community. | F, A |
| **TPR14** | If monitoring is done by a trusted third party (MT3P)[19], all NSPs participating in an SLA must allow the MT3P to retrieve their monitoring data. | F, A |
| **TPR15** | Monitoring data received from other NSPs must be kept secret. | O, F, A |
| **TPR16** | Penalty rules must be accepted by all NSPs participating in an alliance. | A |

### 3.4.3.3. Business Policy Rules

Business policy rules must also be respected by all NSP participating in a federation or an alliance community types.

| BPR# | Description | Type |
|---|---|---|
| **BPR1** | All NSPs must adhere to the community coordination model for conducting business | A |

[19] This is a community decision.

| | | |
|---|---|---|
| | with other community members and also respect the Policy Rules for business interaction with other NSPs outside the community they belong to. | |
| **BPR2** | The NSPs that wish to be part of an ETICS community are obliged to adopt the information propagation and/or management rules of the community and do business transactions as specified by the Business Policy Rules. | F, A |
| **BPR3** | The disclosure of competitive information such as trespassed offers or requests is restricted based on the community-wide regulation. | F, A |
| **BPR4** | It is not acceptable for an NSP to provide access to its end customers with ASQ only via bundled offers unless it also publishes unbundled such ASQ access offers. | A |
| **BPR5** | The way an NSP defines how sessions are admitted and/or bundled in ASQ goods, and how they are transferred among the ASQs already built, is up to its own choice, as long as it does not contradict with possible rules established in the community it belongs to. | O, F, A |
| **BPR6** | Revenue sharing, pricing and SLA penalty schemes of the NSPs should be defined and accepted by all members of the alliance. | A |
| **BPR7** | Cost compensation for providing services, e.g. infrastructure sharing to the ETICS SLAs is a part of the Alliance or Federation agreement. | F, A |

Various topics such as network size, resource availability, capacity, and fault tolerance and how these aspects may also need directions or rules will be considered as various types of ETICS communities are further analysed.

### 3.4.4. ETICS COMMUNITY DEFINITION

An ETICS community is described both in technological and in economics-business terms by a set of network capabilities, lifecycle management and business processes, as well as specific rules for conducting trade and for sharing revenues. The exact specification of these rules will be defined to a large extent by business negotiations in the IC market among the stakeholders. By agreement on common sets of Policy Rules, ETICS communities can be classified as an Open Association, Federation or an Alliance as described above.

An ETICS community could also be formed by homogenous requirements where all NSPs have the same rights and duties. In this scenario, all NSPs could agree to evolve the ETICS community, e.g. from an Open Association to a Federation. In contrary, like e.g. standards bodies, an ETICS community could propose different member levels to its NSPs, thus providing different levels of Policy Rules inside the same community. The ETICS community could be formed around a nucleus of Platinum members which have more duties, but also more rights (i.e., *trust* is the *de facto* standard inside the nucleus) and it could progressively extend to Gold and then Standard members. If the governance of such ETICS community types is more complex compared to a homogenous governance of the community, its members are free to evolve at their individual pace between along different member levels. The figure below shows the position

of 3 types of ETICS community and membership levels, relating Policy Rule constraints to the potential added value in the market place.

The first feedback from carriers[20] suggests first adopting the Open Association model, and then, after observing the market evolution, moving to a Federation or to an Alliance. Currently, few carriers are ready to directly adopt the Alliance model.

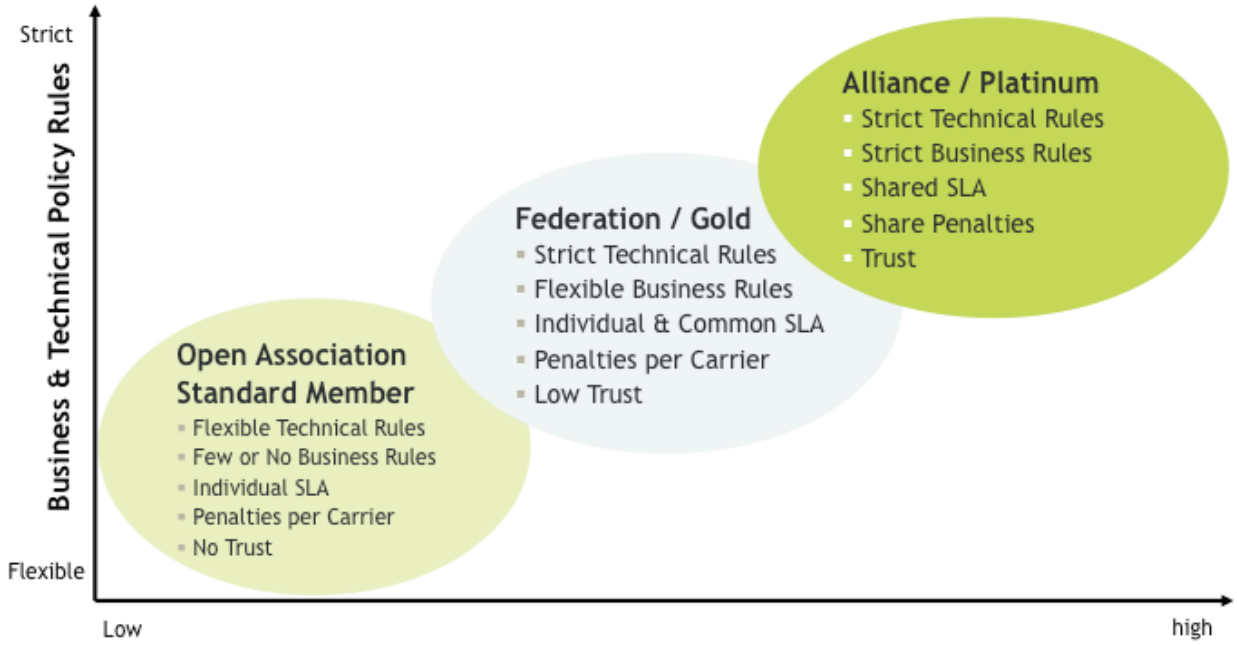


FIGURE 20: ETICS COMMUNITY OPTIONS

[20] Feedback from the 2nd ETICS Carrier Workshop in Berlin, which took place on 20 January 2012.

# 4. ETICS REFERENCE ARCHITECTURE AND SERVICE DEPLOYMENT SCENARIO

This section revises the ETICS architecture defined in ETICS Deliverable D4.2 [ETICS-D4.2]. In a first step, the ETICS global architecture is refined, concerning the identification of the network service and business plane functional blocks and their interactions with both (1) other functional blocks within other NSPs and (2) the network control and data plane.

In a second step, ETICS features and deployment scenarios are recalled from D4.2 and enhanced by the help of UML diagrams. D4.2 has identified different "scenarios" for performing the service composition in order to provide inter-carrier Assured Service Quality (ASQ) paths. These scenarios differ in the used communication mechanism between different ETICS actors, which are required for the service composition. We envision these scenarios to be different features of the same architecture fulfilling the same objective, i.e. providing inter-NSP ASQ paths.

## 4.1. ETICS GLOBAL ARCHITECTURE REVISION

The high-level architecture – as envisioned by D4.2 – serves as starting point for refining the global architecture in this section. Compared with the architecture defined in D4.2, this section goes deeper in identifying the functional blocks within NSPs as well as their interfaces with the neighbouring NSPs. It uses UML modelling for providing a sharper and more precise technical specification of these blocks and their interfaces. After recalling the high-level architecture defined in D4.2 we go in more technical details subsequently.

### 4.1.1. ETICS HIGH-LEVEL ARCHITECTURE

Within the ETICS community, each NSP contributes with its own per-NSP product **offers**. The concatenation of these individual offers represents the end-to-end inter-carrier offer. With respect to this, the ETICS architecture supports two main features: the *on-demand/pull*, and the *pre-computed/push* connectivity offers. In the *on-demand/pull* case, a per-NSP service connectivity offer is provided by an NSP only upon an explicit request for a specific detailed offer. In the *pre-computed/push case*, offers are pre-computed by NSPs "regardless of customer requests". In this case, offers are ready to order. Once an NSP has published an offer, it has the obligation of fulfilling it once it gets an order for it. In an analogy with the clothing industry, the difference between the on-demand and the pre-computed cases can be seen as the difference between made-to-measure and ready-to-ware. Further details are available in [ETICS-D4.2].

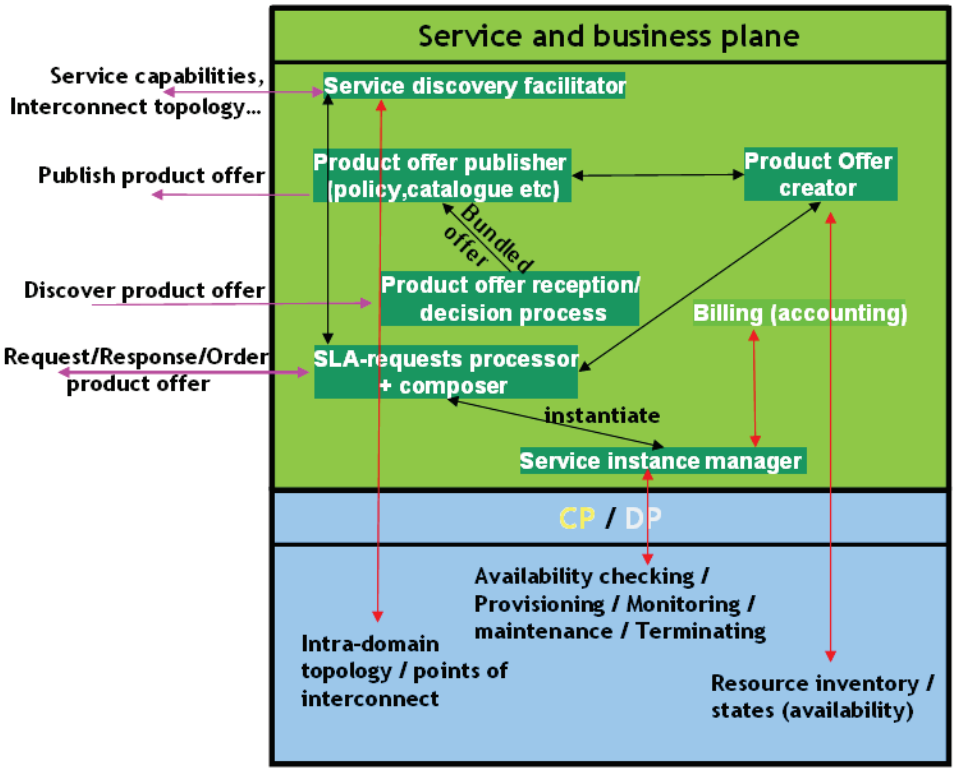


FIGURE 21: ETICS GENERAL REFERENCE ARCHITECTURE (FROM D4.2)

FIGURE 21 shows the ETICS general architecture as defined in the D4.2 deliverable. This high level architecture combines both the pre-computed (push) and the on-demand (pull) axis.

In the pre-computed (push) case, *offers* are first pre-computed and are ready to order. These offers are made available to other actors within the ETICS community. Upon the reception of an order for a given offer, an NSP must be able to provide the offer and instantiate it.

The architectural entities, which are involved in the ***pre-computed (push)*** case, are:

1) The *product offer creator*: it creates (intra) per-NSP connectivity offers.

2) The *product offer publisher*: it is responsible for making these offers accessible to other NSPs or to a centralized facilitator entity (depending on the service composition feature that is used).

3) The *service instance manager*: it instantiates the offers.

4) The *SLA requests processor*: it receives orders for specific offers, and triggers the instantiation process.

Since the offers are available and ready to order, the service composition can be done in this case (depending on the feature, by an NSP or a centralized entity) by computing the best combination of offers that provides an end-to-end ASQ path.

On the contrary, in the *on-demand* model it is not possible to compute directly an end-to-end offer upon receiving a customer request, because offers are not defined in advance. Therefore, a first step is needed in order to find (at least) the list of NSPs that has to be involved to satisfy the end-to-end request. By knowing the list of NSPs that could satisfy an ETICS customer request, each NSP can be requested separately to

provide its offers. This step is performed by using a service discovery facilitator. The facilitator function relies on the exchange of service capabilities to guide the service composition. **Service capabilities** are similar to offers in the way that they specify a technical (sometimes even a business) capability concerning a network connectivity service, e.g. a QoS vector from an entry point to an exit point in the network. However, they differ in the way that they **cannot be ordered directly**; their role is only to facilitate the service discovery.

Therefore, the architectural entities involved in the ***on-demand (pull)*** model are:

1) The *service discovery facilitator*: it relies on the exchange of service capabilities in order to guide the service composition step. This can consist in finding either NSP chain(s) or loose path(s) that can satisfy an end-to-end ASQ path request.

2) The *service requests processor and composer*: it is needed for the exchange of service requests. It also relies on the results of the service discovery function to perform the service composition.

3) The *product offer creator*: it creates intra-NSP offers on-demand.

4) The *service instance manager*: it instantiates offers.

To be precise, the way the different blocks interact with each other depends on the ETICS feature or scenario. More details on this interaction will be provided within the description of each of these features in Section 4.2. The goal of the current section is to start from D4.2 to prepare the identification of the functional blocks as well as the inter-NSP interfaces.

Now, looking at the architecture defined in [ETICS-D4.2], it is possible to identify at least three inter-NSP communication interfaces *at the service and business plane*:

1. A **service discovery interface** that is needed for the exchange of network service capabilities;
2. A **product offer interface** that is needed for the exchange of network service offers;
3. A **service request interface** that is needed for the exchange of service requests between the different network service providers.

We next provide a more detailed version of the ETICS architecture, defining the overall set of functionalities which are to be provided.

### 4.1.2. ETICS FUNCTIONAL ARCHITECTURE

On the basis of the high-level ETICS architecture (cf. Section 4.1.1) this section continues on a more fine-granular level by introducing the ETICS functional architecture. FIGURE 22 depicts the big picture of the ETICS functional architecture, which represents an update of the concepts introduced in D4.2 by detailing some of its functionalities.

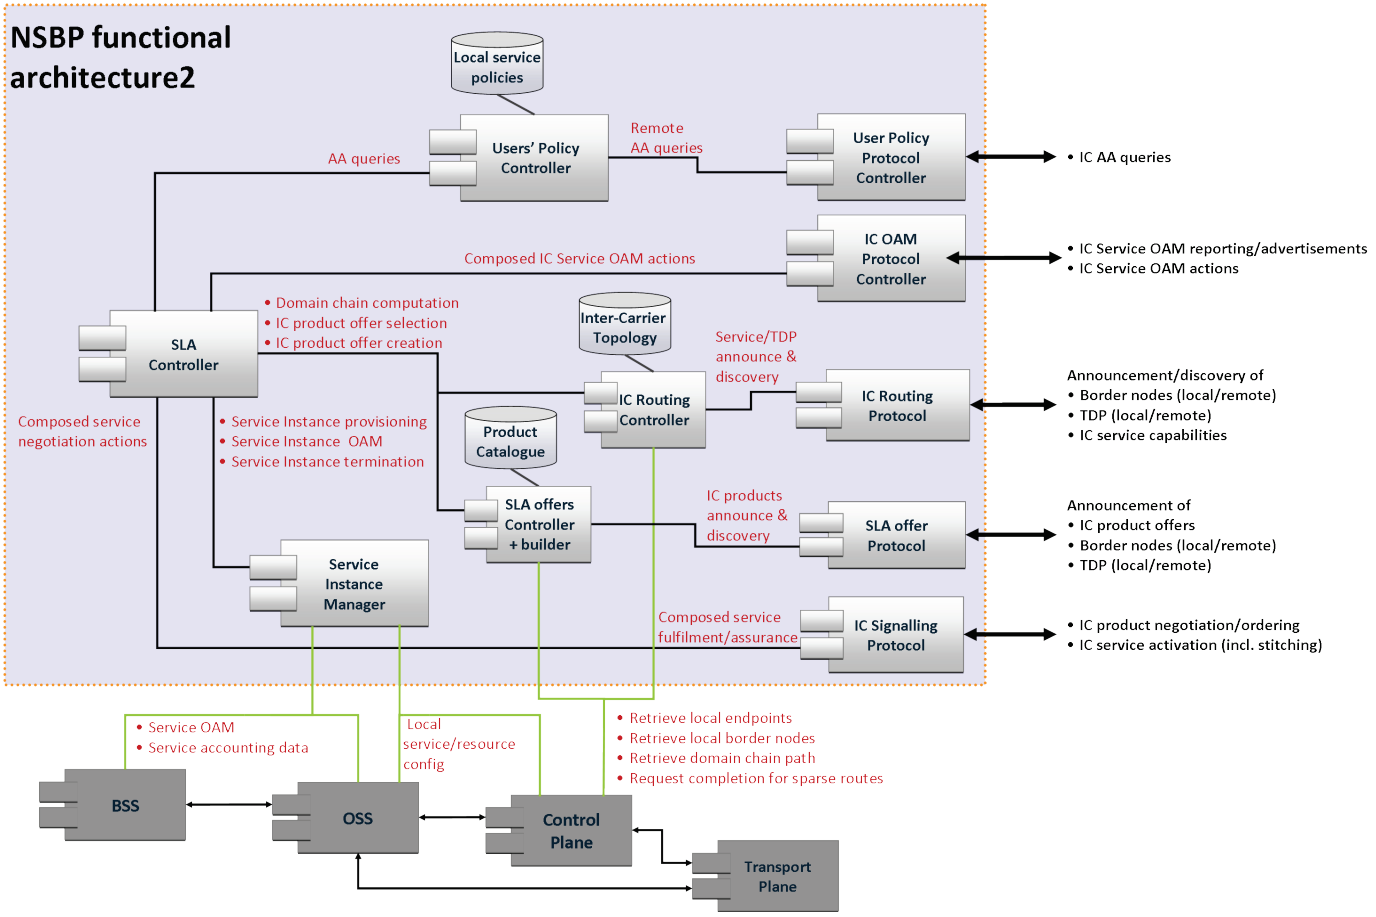


FIGURE 22: ETICS FUNCTIONAL ARCHITECTURE

We have identified three inter-NSP communication interfaces in Section 4.1.1. The first is the ***service discovery interface***. This interface is exposed by three entities in FIGURE 22:

- The *IC routing protocol*: this functional entity refers to the protocol rules and mechanisms by which the service capabilities are exchanged.
- The *IC routing controller*: this controller communicates using the IC routing protocol and is responsible for gathering both intra and inter NSP service capabilities and updating the Inter-carrier topology. This routing controller is particularly important in case of distributed pull as it allows to sending SLA demands to a small set of NSPs relevant for the end-to-end connectivity demand.
- The *Inter-carrier topology*: this entity refers to the repository that contains the inter-carrier topology information or service capabilities.

The second interface is the product offers interface. As in the previous case, this interface is exposed by other three entities in FIGURE 22:

- The *SLA offer protocol*: This functional entity refers to the protocol rules and mechanisms by which the connectivity offers are exchanged.
- The *SLA offers builder*: Similar to the product offer creator in both the pre-computed and the on-demand models, it creates the offers, either on-demand for the pull on-demand scenarios or in

advance in the pre-computed scenario. In the latter case, it updates the Product catalogue database with the local intra-NSP offers.

- The *SLA offers controller*: this controller gathers the inter-NSP offers and feeds the product catalogue database with these inter-NSP offers.

Finally, the third interface is the ***service requests interface***. This interface is exposed by two entities in FIGURE 22:

- The *IC signalling protocol*: This functional entity refers to the protocol responsible for the exchange of service requests between the different NSPs.
- The *SLA controller*: this functional entity refers to the signalling actions related to the composed service fulfilment/assurance.

Other important functionalities depicted in FIGURE 22 are:

- The *Inter-Carrier OAM protocol controller*, which includes mechanisms and protocol messages for the Operation and Monitoring of the composed IC service. It is governed by the SLA controller.
- The *Users' Policy Controller*, which maintains the AuthN/AuthZ data for the ETICS users and services in a single NSP, also implementing the local policy decision point (L-PDP). The policy controller can also interact with other external policy decision points (PDP) or other L-PDP in peering NBSP domains (peer-style) through the Users' Policy Protocol Controller

Adding possibly more interfaces is still under consideration. First, it is important to note that if offers are not enough detailed to allow for (1) end-to-end offer computation and (2) the concatenation and the provisioning of the per-NSP connectivity offers, then a fourth interface might be needed in order to be able to refine the exact path at the network level that needs to be provisioned (similarly, this is valid as well for the service capabilities in case they are not enough detailed). This interface could be at the control plane level though, using one of the already existing techniques like PCE [RFC4655]. Second, depending on the service concatenation technique, a fifth interface might be needed to allow for the concatenation of per-NSP offers (e.g. an overlay RSVP signalling protocol to allow for setting an "interdomain tunnel"). However, this interface could be also part of the IC signalling protocol. These considerations are currently being studied by the consortium.

### 4.1.3. UML DESIGN OF THE ETICS SYSTEM

In order to detail the ETICS architecture, we use an UML model to describe precisely and unambiguously the ETICS system. For that purpose, the following sections gradually introduce the different UML diagrams that have been built. Starting from the Use Case diagrams, continuing with the different class diagrams, we have completed the description with collaborative diagrams as well as sequence diagrams. Even if the UML diagrams could be seen as a simple refactoring of the schemas produced in D4.2 for the ETICS architecture, their principal objectives are to clarify the relation between the different building blocks as well as to describe precisely the architecture minimizing the room for misinterpretation.

#### 4.1.3.1. General Use Case

In the UML modelling, one of the first diagrams corresponds to the description of the actor / role model.

We first remind of the ETICS actor role model (as revisited in Deliverable D4.2) in FIGURE 23 and FIGURE 24. The figures show the actors and the name of the interfaces between these different actors.

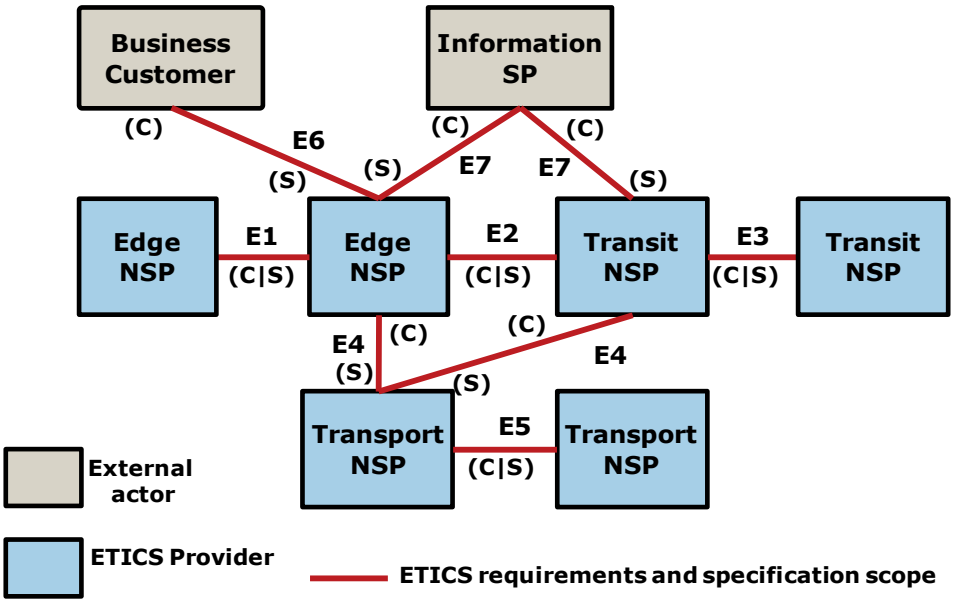


FIGURE 23 ETICS ACTOR ROLES AND INTER-ACTOR REFERENCE POINTS WITH ROLES (C, S)

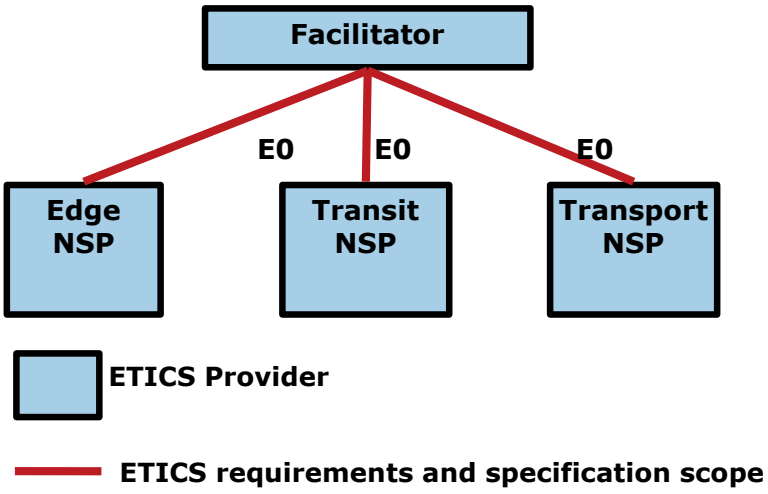


FIGURE 24 FACILITATOR-ETICS REFERENCE POINT

Using the same naming of the interfaces, FIGURE 25 below shows the relation between the ETICS users and the ETICS system.

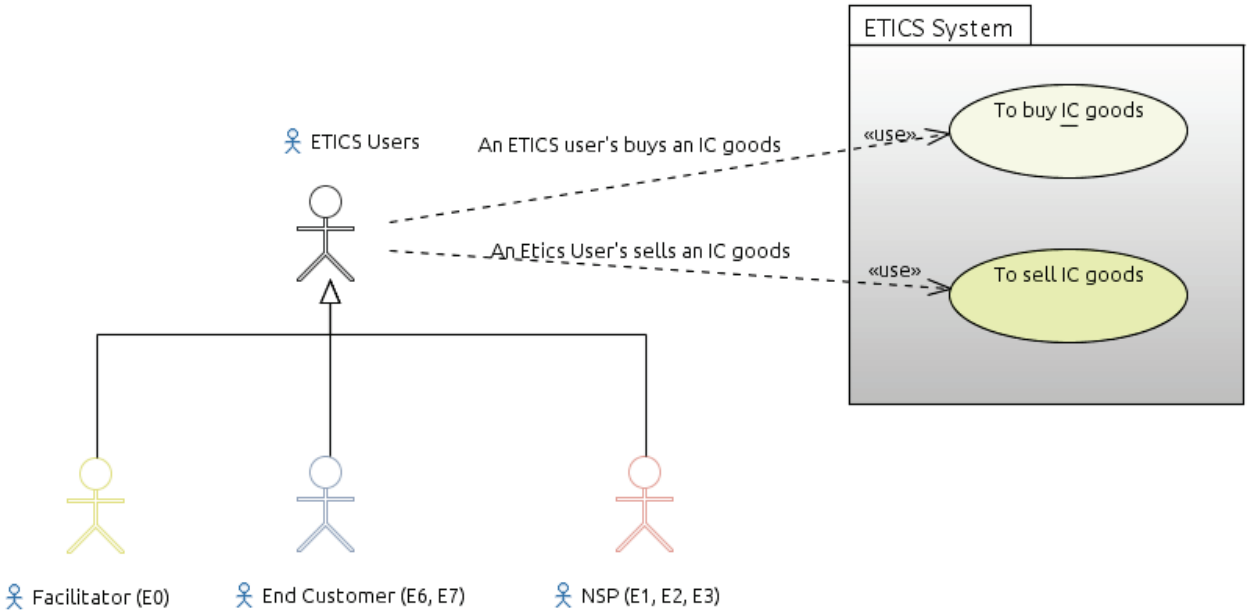


FIGURE 25: UML MODEL OF THE ETICS SYSTEM

The ETICS System offers two very high-level use cases: buying and selling IC goods (cf. FIGURE 25). As a consequence, it seems feasible to reduce the system to simply buying and selling goods, i.e. ASQ goods. Buying and selling IC goods is only available to the ETICS User, which is in fact the abstraction of the three main actors of the ETICS system: the *Facilitator* through the *E0* interface, the *End Customer* (that includes Information Service Provider (InfSP) and Business Customer) through the interfaces *E6*, *E7*, and the *NSP* itself through the interfaces *E1*, *E2*, *E3*. We remind that interfaces E0 through E7 are part of the ETICS actor role model that we (re)defined in deliverable D4.2 [ETICS-D4.2]. The first UML model of FIGURE 25 defines the boundary of the ETICS framework.

#### 4.1.3.2. SLA Life Cycle Use Cases

Nevertheless, FIGURE 25 does not highlight which of the different actors could buy and sell IC goods. In addition, the buying and selling use cases are too generic and do not take into account the different steps in the SLA life cycle, as defined in [ETICS-D4.1] & [ETICS-D4.2. FIGURE 26 below is a more detailed and fine-granular view of the ETICS system use cases, which focuses on the SLA life cycle management.

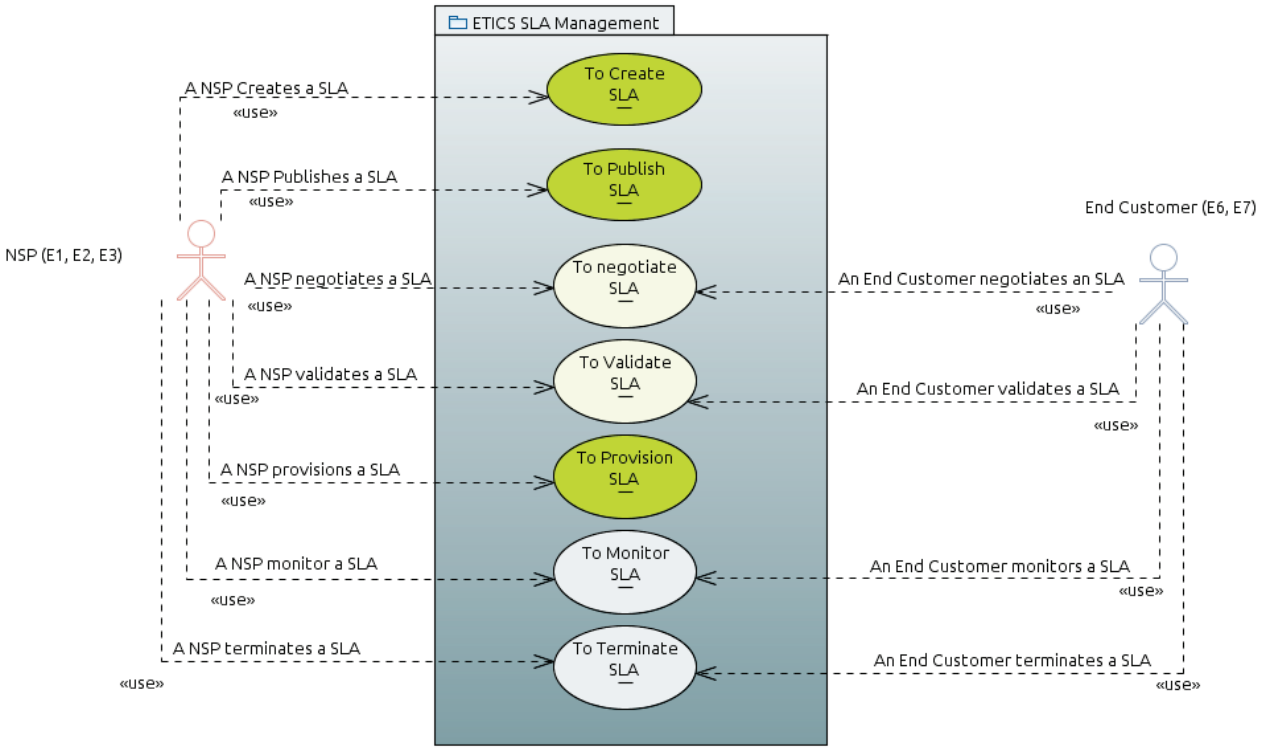


FIGURE 26: UML model of the ETICS SLA Management

This time, the ETICS users have been split in two entities in order to reflect the difference between them. We have the Network Service Provider (NSP) user through the interfaces *E1*, *E2*, *E3* and the *End Customer* through the interfaces *E6*, *E7*. The facilitator (interface *E0*) has not been represented in this UML use case diagram. In fact, the functions provided by the facilitator actor could vary regarding the level of responsibility the NSP would let to him. So, its role could be closer to the End Customer rather than the NSP actor. For the clarity of the UML use case model, we have decided to merely explain that its role could be extrapolated from the different actions described in the UML model.

Of course, we have reproduced the different steps of the SLA life cycle defined in D4.1. But this time, combined with the actors, we are able to better describe it. The End Customer has just access to a subset of the use cases: it could negotiate and validate an SLA that corresponds to the Buying action of the ETICS system. In complement, it could monitor the SLA to verify that the IC goods sold by the NSP are conformant to what it expects, and terminate the SLA when it is not needed anymore. The NSP role is twofold: Buyer and Seller. For the latter one, actions are the same as for the End Customer role: Negotiate, Validate, Monitor, and Terminate an SLA. In the case of the Buyer role, the NSP could create and publish an SLA in order to sell it. Once it has sold an SLA, the last action is to provision it in their network. In addition, the four actions of the seller role (Negotiate, Validate, Monitor, and Terminate) are present but not with the same meaning. The NSP as a buyer provides the action instead of using it.

#### 4.1.3.3. SLA Management Class Diagram

In order to draw the UML Class Diagram, it is important to define the new data type used by the different classes. In our case, we have used those defined in [ETICS-D5.2]. The main advantage of using data types can be summarised in the modularization of the schema and classes, because any change or improvement

in the different parameters of a data type can be automatically reflected in all the containing, referring or derived classes.

The class diagram depicted in FIGURE 27 below illustrates the ETICS SLA management system. It consists of four main packages directly related to the SLA:

- **SLA Manager** groups all classes involved in the management of the SLA that occurs during the Negotiation, Validation, Termination, and Service Assurance steps of the SLA life cycle. A dedicated class is used to communicate with the rest of the ETICS system during the aforementioned actions.
- **SLA Offers** groups all classes that belong to the creation, certification and publication of SLA offer actions. It also embeds the class that communicates with the rest of the ETICS system to exchange SLA offers. This is the second external interface of the system. This package is in charge of the creation and publication actions in the PUSH model only.
- **IC Routing** groups all classes devoted to the Inter-carrier Routing behaviour. In particular it defines the way Network Capabilities are built against Network Topology as well as how this information is exchanged with neighbouring NSP(s). This package is in charge of the creation and publication actions in the PULL model only.
- **Business & Policy** groups all classes involved in the business and policy rules that govern not only how the SLA offers and Network Capabilities are built, but also how the service composition shall be done in the SLA Management package. “Billing and Accounting” completes the package.

Two other packages are indirectly linked to the SLA management. These packages have a lower detail level. Indeed, those are not really part of the SLA Management, but due to their strong interaction with it, it is important to model them in the class diagram:

- **Network package** aims to group all tasks related to the real network. At least Monitoring, Measurement, Service instantiation, Topology acquisition, and Network Configuration are part of this package.
- **ETICS UI** groups all features related to the User Interface that the system exposes to its customer. It models the interface *E6* and *E7* disregarding the protocol, interface, or API exposed by this package. The purpose is to show that SLA composition could be triggered by both, another SLA Manager instance (from a neighbouring ETICS system through *E1*, *E2*, or *E3* interfaces) or by an ETICS users through the *E6* or *E7* interfaces.

#### 4.1.3.4. SLA Offers Package

It is composed of 3 classes: SLA Offers Builder, SLA Offers Controller, and SLA Offers Protocol. The Builder is in charge of composing the offers. To this end, it relies on information from the Network Package: topologies and measurements. With all this information, it starts to pre-fill SLA templates and passes them to the SLA controller, which refines the SLA template with the rules acquired from the Business & Policy package. In this way, the own SLA offers are stored in the catalogue. Based on the policy rules, the SLA controller sends offers to the SLA Offers protocol class that must be published to the other NSPs. In turn, it gets the offers received from the neighbour NSP. If the option is supported, the SLA controller could re-bundle the received offers in order to enrich its offers catalogue. Of course, foreign offers are also stored in

the SLA catalogue. Finally, the SLA Offers controller exposes functions to the SLA controller, in particular the "Get SLA Offers" action, when this last one performs the SLA Service Composition using the Push model. Note that SLA offer actions are running in background. Compared to the SLA Management package actions, they are running in parallel.

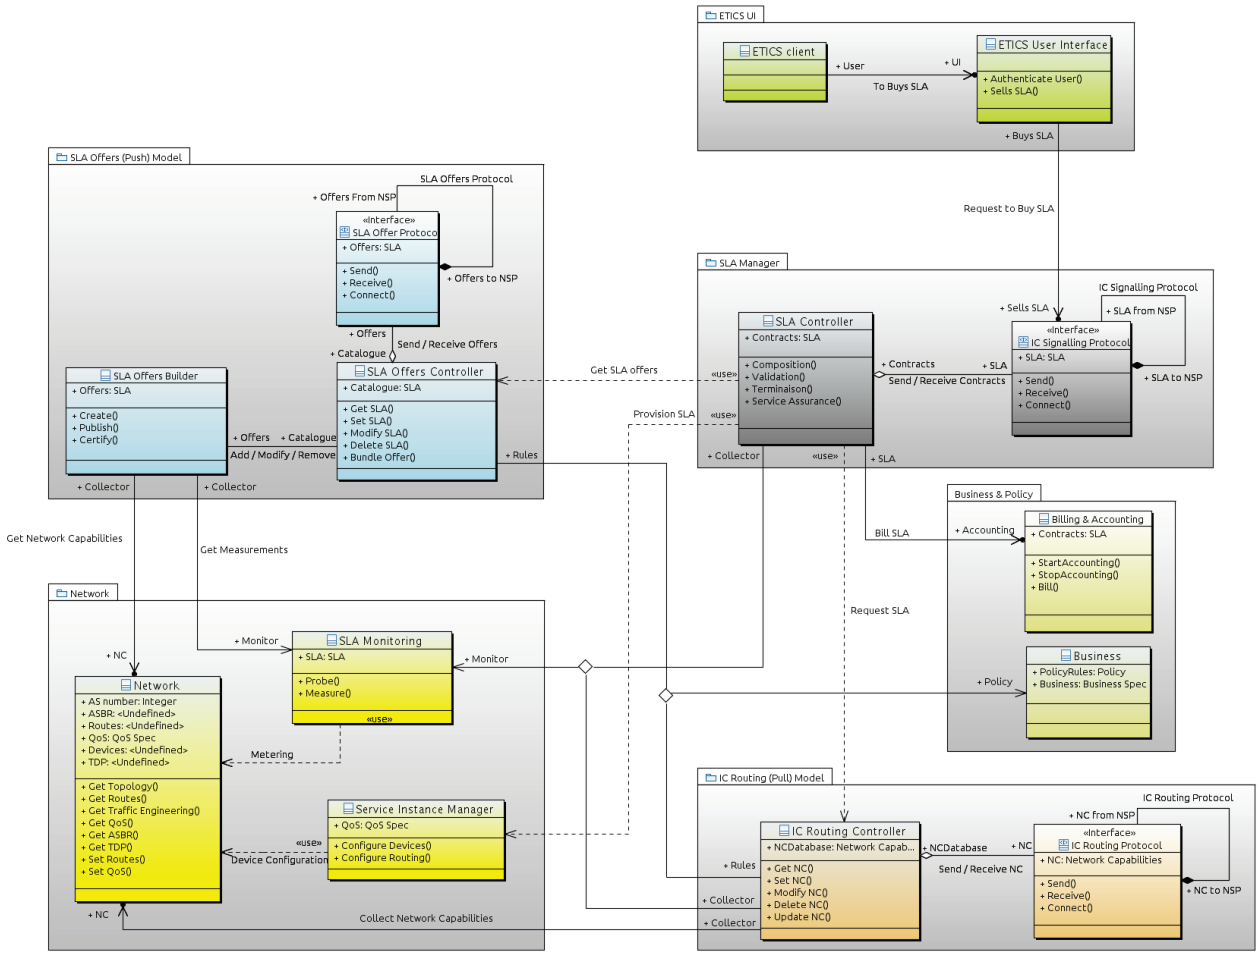


FIGURE 27: ETICS SLA Management UML Classes Diagram

### 4.1.3.5. IC Routing Package

The package embeds only two classes: the IC Routing Controller and the IC Routing Protocol. The latter one is used to exchange Network Capabilities with the rest of the ETICS systems. The IC Routing Controller is in charge of computing the Network Capabilities for its own network and of publishing them through the IC Routing Protocol. To compute the Network Capabilities, the IC Routing Controller collects topology information from the Network package and complements it with measurement information. Then, it applies rules from the Business & Policy package before publishing the Network Capabilities. All Network Capabilities (its own ones and those coming from the other NSPs) are stored in a database that allows the IC Routing Controller to reconstitute the topology at the AS level (i.e. at a upper hierarchy view rather than the intra-domain standard topology; please refer to Section 6.1 in [ETICS-D5.2] about the Hierarchical Traffic Engineering description). Finally, the IC Routing Controller exposes functions to the SLA controller, in

particular the “Request SLA” action. It is used within the SLA Service Composition process when the Pull model is performed.

Note that IC Routing actions are running in background. With respect to the SLA Management package actions, they run in parallel.

#### 4.1.3.6. SLA Management Package

This package groups all classes, which are involved in the composition, validation, termination, and monitoring of the SLA. In particular, the package manipulates contracts compared to offers for the SLA Offer package. It is composed by the IC Signalling Protocol in charge of exchanging information during the different steps of the SLA life cycle, which are: composition, validation, and termination. The main class is the SLA Controller that embeds the composition, validation, termination, and monitoring function. For that purpose, the SLA controller will use the SLA Offer Package when working in PUSH model and the IC Routing Package when using the PULL model. Requests to compose a new SLA contract come from the ETICS UI or from another NSP through the IC Signalling Protocol class. The SLA controller has also interaction with the Business & Policy Package for the Accounting and Billing process and with the Monitoring & Measurement Package to request monitoring of a given SLA contract. It maintains a database where all SLA contracts are stored.

#### 4.1.3.7. Business & Policy Package

This package is twofold: (1) It provides Business & Policy rules to the other packages in order (for the NSP) to control how the information (SLA offers, Network Capabilities, SLA contract) is computed locally and exchanged with the other NSPs. The rules are stored in a database filled by the operator through a user interface (GUI or Web page). (2) The second part corresponds to accounting and billing. Each time a new SLA is contracted, an accounting ticket is issued in order to start the billing process and to bill the ETICS client. When the contract is finished, another accounting ticket is issued to stop the billing process. This package is more or less standard. Hence, we aim at reusing existing software packages for that purpose.

#### 4.1.3.8. Network Package

This package describes all information related to the Network of the NSP, but without going into too much detail. The goal is to describe more the interaction between the network and the other SLA packages rather than model the network in UML. Nevertheless, we have made the distinction between (1) the Network itself, which provides information on topology, QoS, routing, etc. to the SLA Offer and IC Routing packages, (2) the Measurement & Monitoring system, and (3) the Service Instantiation. The latter is in charge of configuring the device in order to enforce the route and the QoS into the Network. Again, as it depends on the choice of NSPs and on the technology used, we keep this class generic. It could be refined later per technology, e.g. for MPLS-TE. The Measurement & Monitoring class aims to provide complementary information for the SLA Offers and IC Routing packages in order to build Offers and Network Capabilities. Finally, the Measurement & Monitoring could be triggered by the SLA Management package in order to monitor a given contract.

#### 4.1.3.9. ETICS UI Package

This package shows how an ETICS Users could buy an SLA contract through the ETICS User Interface that represents the *E6* and *E7* interfaces of the generic ETICS architecture.

### 4.1.4. MONITORING ARCHITECTURE

In previous deliverables, network monitoring has been identified as important building block of the overall ETICS ecosystem. It has been developed in architectural deliverables D4.1 (ETICS-D4.1) and D4.2 [ETICS-D4.2], as well as in design, specification, and implementation related deliverables – i.e. D5.2 [ETICS-D5.2] and D5.3 [ETICS-D5.3]. Moreover, the topic has been addressed by the test-bed design (D6.1 [D6.1] Section 2.7) where a network topology suitable to check aspects of the ETICS network monitoring has been presented.

In this section, the term "domain" is used with respect to network domains where "domain" means an adjacent part of the (Inter-)network infrastructure that is under the administrative control of one authority.

*Note:* For technical reasons, non-adjacent parts of the network infrastructure must be considered as different domains even if they were under administrative control of the same authority.

In the next few paragraphs, we describe which aspects of network monitoring in ETICS have already been addressed in previous deliverables. However, starting with Section 4.1.4.1, network monitoring is described in a self-contained way such that there is no need to check external references, even though they are provided throughout the text.

Network Monitoring has been discussed so far in [ETICS-D4.1] (D4.1 Sections 3.2.6 Monitoring & Measurement, D4.1 Section 5.3 End-to-End SLA and Network Monitoring, Annex C Network Monitoring) and in [ETICS-D4.2] (D4.2 Section 6: SLA Monitoring and Assurance Architecture). The latter section examines the ETICS monitoring system and related interfaces to the Control, Data and Management Planes, and thus it represents the main source of information for the present deliverable.

Furthermore, a first attempt to identify the interface points of the monitoring subsystem with the core ETICS system has been performed in D5.2 [ETICS-D5.2] in Section 7.2.

#### 4.1.4.1. Purpose of the ETICS monitoring system

The purpose of network monitoring in ETICS is threefold:

(1) Verification of SLA fulfilment,

(2) Debugging ETICS installations and configurations, and

(3) Educational purposes.

**Regarding (1)**: Business analyses did not yet yield definitive information as to whether the ETICS business cases call for fully fledged monitoring system like NMON, or if other, less technical measures, will represent the method of choice. However, since QoS-enabled transport services ("goods") are envisioned to be priced more aggressively, it is assumed that disputes between domain owners along the data path are hardly solvable without monitoring technology. Having said that, it is obvious that the monitoring solution must be carefully designed in order to keep monitoring costs low. This proposal seeks to take this into account. Moreover, NMON can serve as an input to estimate customer's Quality of Experience (QoE), and consequently countermeasures can be taken ***before*** the QoE reaches a lower threshold, which helps to avoid user churn due to bad QoE. Mapping of QoS parameters (delay, jitter, etc.) and multimedia parameters (codec, motion level, etc.) into QoE levels, that is to say service quality as perceived by the end

user, is indeed part of the ETICS framework. QoS parameters, which are the output of the NMON architecture, could provide input to the QoS-to-QoE mapping box.

**Regarding (2)**: Especially during the set-up process of QoS enabled services, NMON can be used to precisely spot the points of failure, which results in faster time-to-market and lower initial costs.

**Regarding (3)**: In the case where over-provisioning is not possible or economically meaningful, traffic behaviour is hard – if not impossible – to predict. NMON in this respect can serve as enabler for the research area of inter-domain QoS, which still has encompasses a large research potential. Related work has already been studied and described in the [ETICS-D4.1], Annex A (State of the Art), which will not be replicated here. Please see the mentioned chapter for further details.

A more detailed description of the NMON purposes and scenarios is available in the D4.1 Section 5.3; furthermore, a more detailed description of NMON terminology and methodologies can be found in the Annex C of the same deliverable.

#### 4.1.4.2. Metrics with high relevance for ETICS network monitoring

For several years, the measurement community has been defining metrics in order to perform measurement in networks, in particular in ITU-T (International Telecommunication Union Telecommunication Standardization Sector) and in the IPPM (IP Performance Metrics) working group of the IETF. The metrics can be classified into three categories:

- Local metrics: performed at the destination, they show mainly the variation of the service received.
- One-way end-to-end metrics: relevant for evaluating the network service or for performing measurements for non interactive applications.
- Two-way end-to-end metrics: relevant for evaluating interactive applications or TCP-based transport services.

Within ETICS we will only focus on the main metrics relevant for evaluating the QoS (delay, jitter, bandwidth and loss) experienced in the network, either targeting the full end-to-end path or single interconnect links:

**Local metrics**

We will focus on the jitter, also called variation in packet delay (RFC 3393). When monitoring the network service, this metric gives an indication of the network state evolution (getting into congestion or not) since it shows the dynamics of queues in the network. Additional metrics such as packet reordering (RFC 4737) (for example, applications using the TCP transport protocol) may be considered.

**One way end-to-end metrics**

- One-way delay [RFC2679]: this is the transfer delay experienced by packets in the network. This is in general the most important metric when estimating the network service for non interactive or non TCP-based applications since the symmetry of the network paths cannot be assumed.
- One-way loss [RFC2680]: this is the loss experienced by packets in the network and complements naturally the information given by the one-way delay.

- Available bandwidth [RFC5136]: this is the amount of unused bandwidth on a link.
- Obtained bandwidth: this is the throughput actually consumed by a given application.

**Two way end-to-end metrics**

The two-way delay, also called round-trip time (RFC 2681) is relevant when the return path has to be considered for the communication.

In order to evaluate the QoS obtained from the network, end-to-end monitoring of the relevant metrics is a priority for ETICS. On the one side, one-way metrics are more appropriate since they do not assume symmetry of the network (services), on the other side, they are noticeably more difficult to obtain since they usually require synchronization of probes and/or end hosts on both "ends" of the data path and always require some form of correlation of measurements performed at both ends.

#### 4.1.4.3. Monitoring approaches

In the following subsections, various network monitoring approaches are presented.

"OAM monitoring" and "autonomous monitoring" are both used in the ETICS solution in parallel, as they are orthogonal to each other, which will be explained in these sections.

"Centralised monitoring" and "Active flow based monitoring" have also been evaluated and are for further study (i.e. they are currently not foreseen for use in the ETICS solution).

##### *4.1.4.3.1. OAM monitoring*

OAM (Operations, Administration, and Maintenance) is a general term used to describe the processes, activities, tools, standards, etc., involved with operations, administration, and maintenance activities mainly used in the context of computer networks. A set of standards, e.g. ITU-T Y.1731 [Y1731] and [IETF-DR-4], to name only some prominent ones, has already been defined. In the ETICS context we restrict ourselves to the OAM aspects of monitoring – i.e. determining QoS metrics – of computer networks.

This topic had been introduced later in the ETICS project and has not yet been described in previous architectural deliverables.

OAM standards, operating on Layer 2 of the OSI reference model [ISO-OSI], are already implemented in most of recent network equipment and can readily be used for Layer 2 (L2) monitoring purposes., e.g. Link Trace and Ethernet Loopback.

However, besides these advantages, OAM has some drawbacks: firstly, it suffers from scalability issues, and secondly, it lacks security mechanisms needed for an inter-operator use. For this reason, the OAM monitoring is used only for intra-domain purposes.

OAM concepts will be applied for the fast reaction to possible faults and not with specific quality concepts. In a data network there are many points where OAM could be applied – Figure 28 below is highlighting some of them:

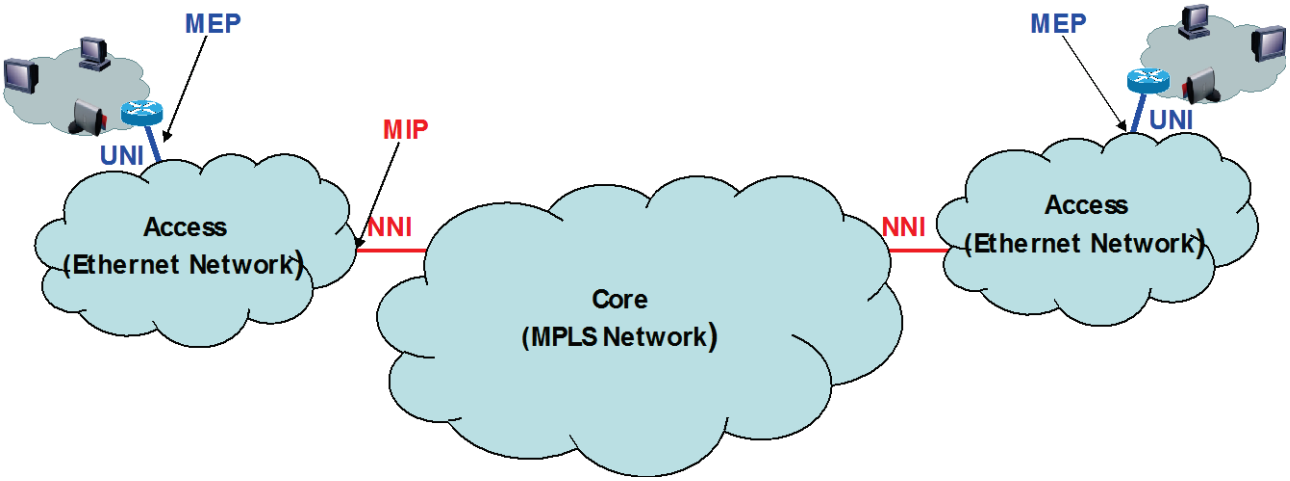


FIGURE 28: OAM MAINTENANCE ASSOCIATION END POINT/ASSOCIATION POINT

By controlling some specific point in the network the status of a connection is managed. These points are divided in two main categories: MEP and MIP.

- MEPs (Maintenance Association End Point) are the extremities of the connection that are under control. The data retrieved are important for verify the quality of the connection itself.
- MIPs (Maintenance Association Intermediate Point) are other points under control. These are optionally defined, and it increases the speed of the fault detection.

The network example depicted in the above figure focuses on the **Association Points** definition that, with the **Trustiness,** can be considered one of the main drawbacks of the OAM functionality.

- **Trustiness:**

  The standard bibliography makes no distinction between inter- and intra- domain scenarios.

  However, the **intra**-domain case has only limited complexity, because a single operator controls the whole network. This strongly differs in the **inter-**domain case, where the path crosses networks of different administrative controllers. In theory, each operator has to show its internal network structure to the “community” in order to provide a common policy control.

- **Association Point Number:**

  The definition and distribution of the Association Points are directly controlled by the operator, which possesses the complete topology of its own network. Since every active point has to react "quickly" once it is solicited, a huge amount of points (mainly MIPs, but also MEPs can be in the context) must be carefully considered.

#### *4.1.4.3.2. Centralized monitoring methodology*

A first solution is to be considered in which an authorized trusted third party is in charge of managing the monitoring system, i.e. a centralized monitoring mechanism as depicted in FIGURE 29. This is a very strong assumption which implies that each autonomous system allows this third party to administrate its monitoring equipment. Although this simplifies the gathering and the correlation of measurements, it assumes that there is communication and message exchanges in place between the central entity and the measurement points of the individual autonomous systems. Further, it implies a high level of trust between the autonomous systems and the third party. Many issues are being raised by such a solution since the autonomous systems are obviously sensitive with respect to privacy issues and as they will therefore not easily provide access to the monitoring probes. We believe that this solution is unlikely to be deployed, except in the case where the autonomous systems are under the administrative control of the same authority.

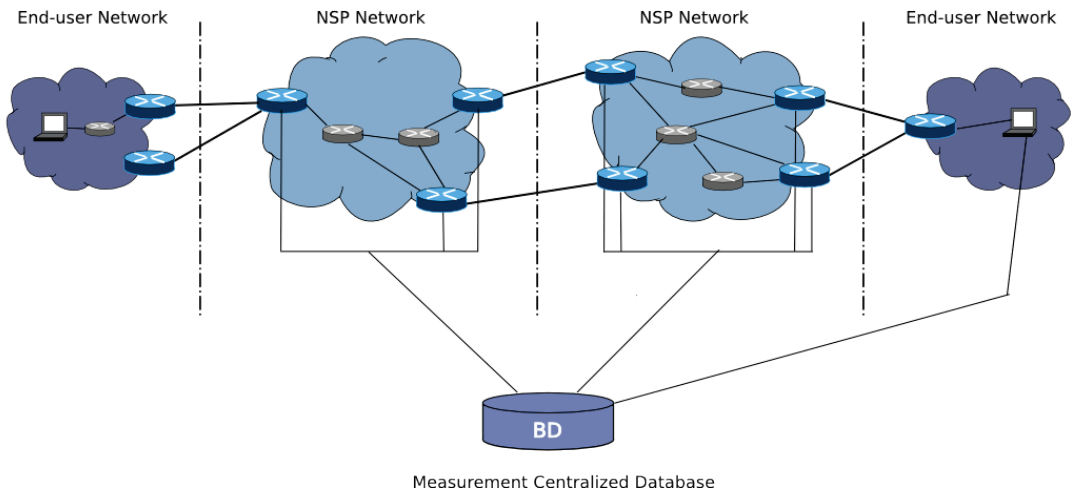


FIGURE 29: CENTRALIZED MONITORING ARCHITECTURE

#### *4.1.4.3.3. Active flow based monitoring methodology*

Active flow based measurement architecture may also represent a potential solution (see Figure 30). This architecture simplifies the issue of correlating the domain-local packet measurements to infer QoS metrics experienced by the measured packet. It addresses the storage issue but it also has got some drawbacks. With active measurements, the performance experienced by a probe flow provides a good approximation of the QoS experienced by an application. The precision of the measurement depends on the aggregation point, the synchronization of measurement equipment and the intrusion of the probing flow. The metrics evaluated are for example: one way delay (OWD), one way losses, etc. The principle is to have equipment

with measurement functionality in the path, which is able to add a timestamp to a probe packet in order to compute the OWD by subtracting the ingress timestamp from the egress timestamp, the former of which is stored in the probe packet. This architecture addresses the correlation issues of the measurement, but it assumes that the border equipment of the NSPs is: (1) able to recognize the active probe packet, and (2) able to timestamp those packets. Effectively, this translates to an update of those nodes, which seems difficult to achieve in short time. For this reason, this architecture is interesting but appears to be hardly deployable for the moment.

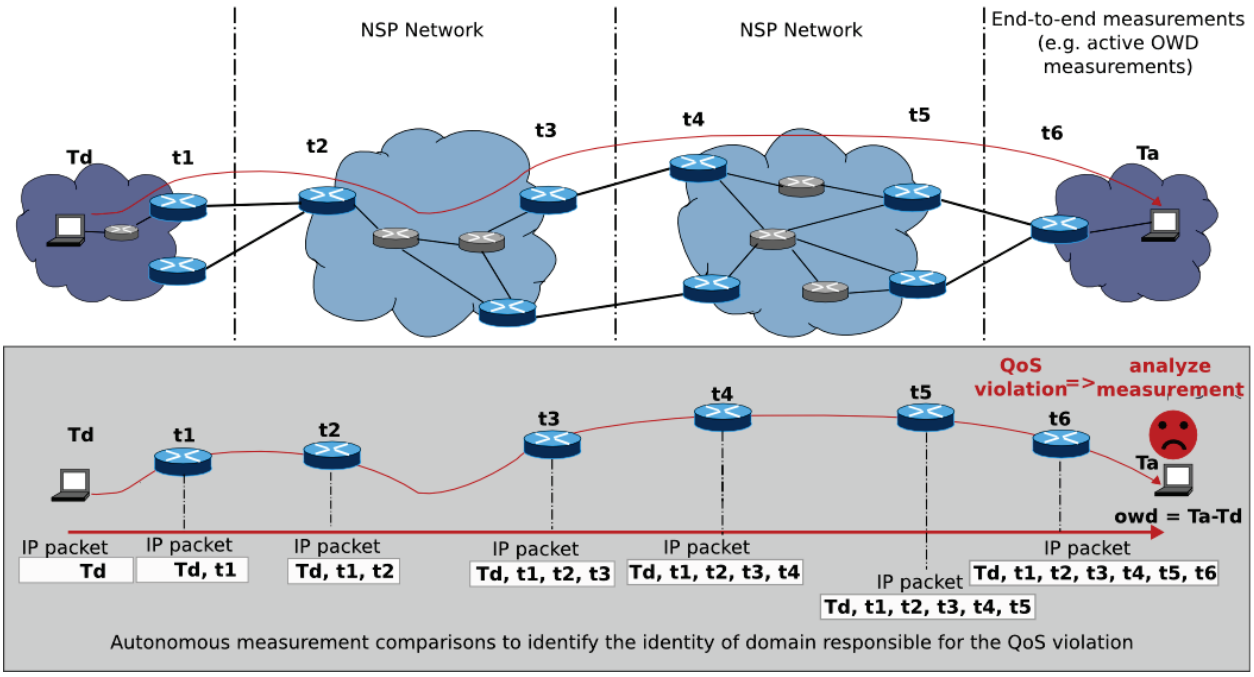


FIGURE 30: ACTIVE FLOW BASED ARCHITECTURE

#### *4.1.4.3.4. Autonomous monitoring methodology*

Another option is to consider that each autonomous system will administrate its own monitoring architecture independently of the other systems. In this autonomous monitoring architecture, SLA violations on the end-to-end path will be detected by end-to-end active or passive measurements (cf. Figure 31). Active measurements would usually be launched by the source end-user network (or by local measurements at the destination end user network, depending on the type of metrics).

Detection of SLA violations will set off a mechanism to access monitoring information at the autonomous systems (domain) level in order to locate QoS degradations (see Figure 32). The access to the monitoring data will be provided by the NSPs. They may send their measures in response to a solicitation from the monitoring system or authorize other domains to access the data. This latter solution asks for a hierarchical architecture which allows finding and transferring only this data, which is needed for a certain request/measurement. This is necessary because a full transfer of all monitoring data would produce way too much traffic (cf. [ETICS-D4.2] Section 6.4).

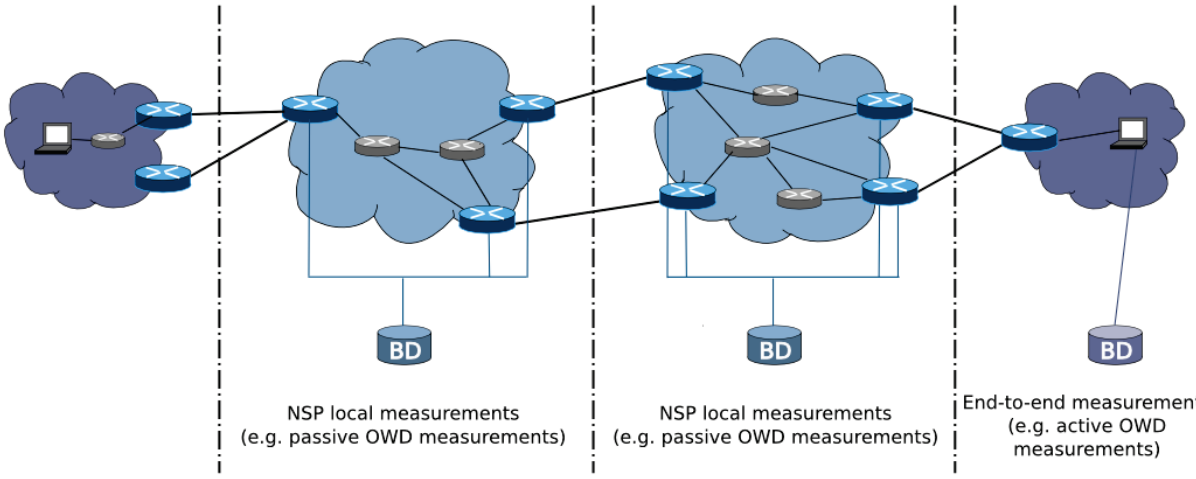


FIGURE 31: GENERIC AUTONOMOUS MONITORING ARCHITECTURE

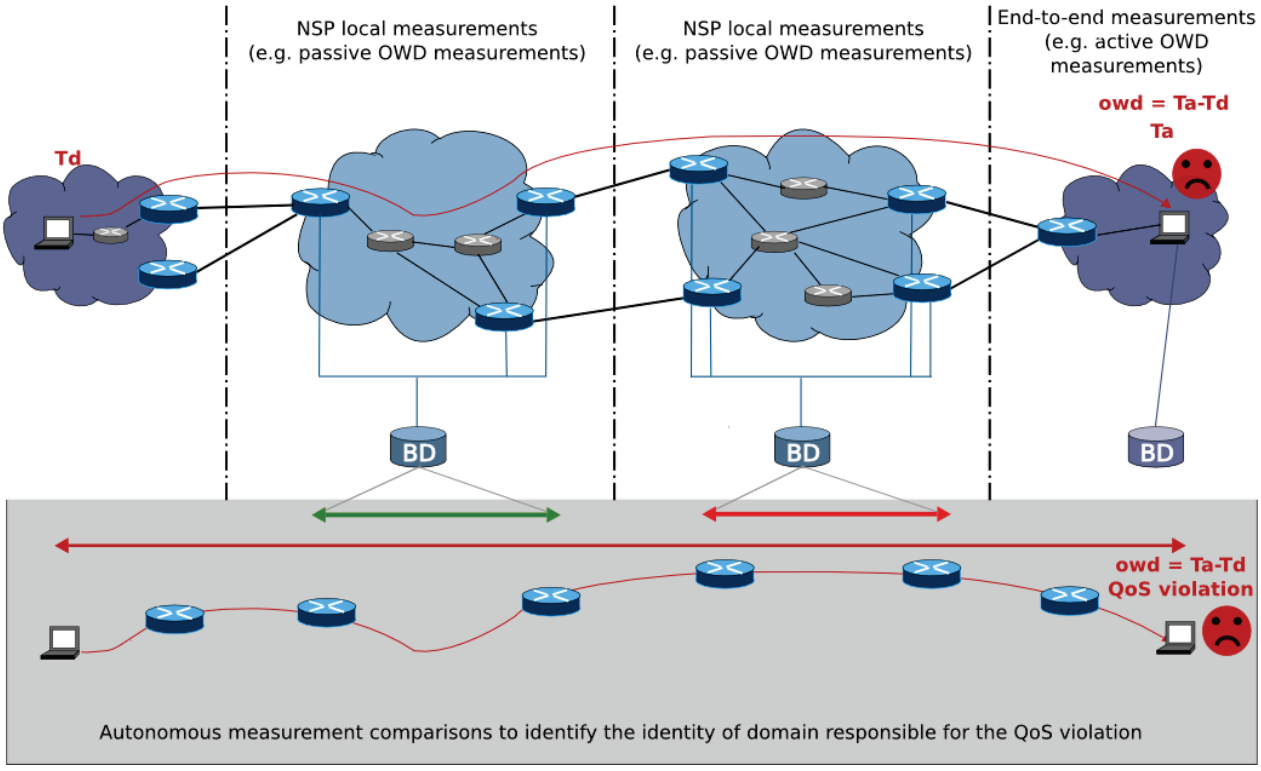


FIGURE 32: QOS DEGRADATION LOCATION IN CASE OF A SLA VIOLATION IN THE AUTONOMOUS MONITORING ARCHITECTURE

The hierarchical monitoring architecture utilising this autonomous monitoring methodology is called “NMON” and is described as follows:

Functionality of the NMON system must carefully be balanced against cost thereof, so the overall monitoring system may not be too complex. This hierarchical approach described here is simple and straightforward and fits well to the autonomous monitoring methodology. Moreover, scalability issues and security/confidentiality concerns can also easily be addressed. This approach is meant to implement the ETICS monitoring system quickly and with reasonable effort.

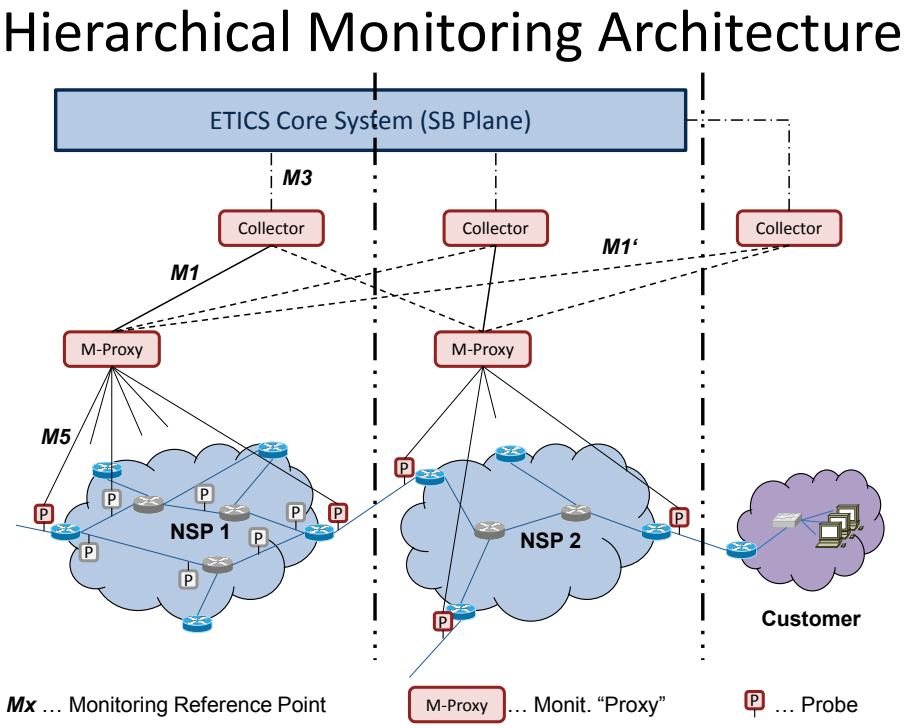


FIGURE 33: HIERARCHICAL MONITORING ARCHITECTURE

FIGURE 33 shows the basic functional entities of the hierarchical approach. The probes are deployed in the network (physical layer) at least on any ICI/TDP (cf. section 3.2) of any (monitored) link in both directions. An operator may also decide to use the system for monitoring the network internally and to that end deploy probes also within the network, which is demonstrated for *NSP1* in FIGURE 33 (greyish probe symbols).

The M-Proxy (Monitoring-Proxy) function acts as point-of-contact for the other monitoring functions. For availability and/or load-balancing purposes, more than one instance of the M-Proxy function could be deployed per domain. Other functions, such as caches and databases for storing data where applicable (e.g. link-based performance data), could be co-located with the M-Proxy function.

The collector function is part of the front-end and is used to collect, correlate, and evaluate the data retrieved from the probes. The collector function also needs to talk to the ETICS ("non-monitoring") system (via M3) to retrieve necessary data such as routes of established ASQ paths and SLA/SLS information. None, one, or more instances of the collector function can be deployed to either

- NSP domains, including the end customer's domain, or,
- A number of trusted third parties (monitoring service provider).

Which scenario shall be used for the final ETICS architecture is subject to further studies, but is certainly influenced by the extent trust relationships are established between actors. However, the decision where to deploy instances of the collector function does not change the reference points.

*M1* is the reference point between the Collector and the M-Proxy function within the same domain. *M1'* is basically the same reference point as *M1*, but improved by some authentication / encryption mechanisms

needed for inter-domain, untrusted relations. *M1/M1'* protocols should support forwarding and redirection of requests/connections for availability and load-balancing reasons.

*M5* is the reference point between a probe and the M-Proxy function, which is always within the same domain.

Besides *M3*, the main function of which has been described previously, the main function of reference points *M1*, *M1'*, and *M5* is to transport information for configuring the probe and to transfer monitoring data from the probe towards the collector function. For all reference points shown, functions and protocols will be specified in WP5.

NMON covers two mechanisms: active and passive monitoring, both of which are developed on top of the same architecture as shown in FIGURE 33. Active- and passive monitoring have in common that they both work on OSI layer 3 (in contrast to OAM monitoring) and both use hashing in some way. This means that packets are captured at different points on the path and are correlated afterwards by means of hashes. To this end, monitoring information is retrieved by the collector function through monitoring proxies (M-proxy in FIGURE 33).

The difference of active and passive NMON is that in the former case packets are inserted for the sole purpose of getting QoS values, while in the latter case only packets of the user's application traffic are captured without injecting additional packets into the network.

**Active NMON:**

Similarly to the Active Flow Based methodology, active NMON is based on active probe packets (i.e. packets injected into the network), which are inserted at the ingress edge and sent through the network towards the egress edge for the sake of monitoring the QoS that they receive. QoS parameters that can be monitored by active mechanisms include, but are not limited to, end-to-end availability, delay and available bandwidth.

At the ingress edge, the emission timestamp is embedded in the payload of active probe packets; at egress edge the reception timestamp is logged. End-to-end delay can be computed at egress edge or at a collector level as the difference between the emission timestamp and the reception timestamp.

One-way latencies measurement has been standardized by OWAMP (One-Way Active Measurement Protocol [RFC4656]). Available bandwidth can be obtained by post-processing the one-way latencies of a series of active probe packets. Different methods exist for doing so, an example of those methods being *Forecaster* [NeHa]. End-to-end availability is a straightforward by-product of active one-way latency monitoring, as any packet which is not received at the egress node is logged as dropped.

In ETICS one-way metrics (availability, latencies, available bandwidth) must be monitored from ingress node to egress node but also per NSP since it is important to locate the faulty NSP in case of SLA violation. As in the case of Active Flow based monitoring, active probe packets are also marked at the ingress node such that they can be identified and evaluated by equipment on the path. To this end, a "magical value" and a sequence number are inserted into the payload of probe packets. The magical value identifies a monitoring job, and the sequence numbers identify the successive packets in the train of active probe packets that correspond to a monitoring job.

Observe from this discussion that there is a lot in common between Active NMON and Active flow based monitoring. The difference between them stems from the behaviour of intermediate equipment on the path. In the case of active flow based monitoring, the intermediate pieces of equipment are active. They includes a timestamp (i.e. the timestamp of that packet at intermediate equipment) into the payload of active probe packets. At the egress node an active probe packet payload consequently carries a sequence of timestamps, and the contribution of each NSP to the global latency can be computed as the difference between two of those successive timestamps. On the contrary, in Active NMON all intermediate pieces of equipment are passive. They recognize active probe packets by a series of features among which are the magical value and sequence number (mandatory), as well as port number and protocol (optional). They log the timestamps of active probe packets into a memory efficient data structure that we call sketch and which is based on different hash functions. The timestamp of active probe packets can be retrieved by a collector that is able to send monitoring requests to ingress nodes, egress nodes and intermediate check points through the monitoring proxies.

**Passive NMON:**

The ETICS passive monitoring system is also based on hashing: for every packet at an observation point (cf. FIGURE 33) a hash value is calculated from the packet headers (fields which might change *en route* must be masked out) and part of the payload of the packet. The same hash function must be applied at all observation points. Consequently, if one packet passes several observation points, it is always gets assigned the same hash value. In this way, the packet can be "tracked" on its way through the network by finding the same hash value at different observation points. The hash value, along with the time-stamp when the respective packet was captured, is stored directly on the probe or on a fast storage which is attached using a high-speed I/O interface.

In modern networks with transfer rates of 10GBit/s and above, it is not feasible to store all hash ⇔ timestamp values, because this would consume too much storage space. Therefore, a filter based on the hash value will be put in place. Of course, in order for this to work, the same filter must be applied at all observation points. In this way, it is possible to store only a portion of the {hash, timestamp} pairs in order to reduce storage space and I/O demand.

The actual calculation of traffic metrics is done by the Collector function (cf. FIGURE 33), which retrieves the timestamps from the probes via the M-Proxy function. The M-Proxy primarily acts as contact point for all monitoring data requests. After verifying access rights, the Collector might be redirected directly to the probe or storage where to retrieve the data from in order to shorten the path the monitoring data takes (this is not shown in FIGURE 33). Furthermore, the M-Proxy function can be deployed several times in a network for load balancing and redundancy reasons.

Please note that in FIGURE 33, the Collector function is shown in the operator's networks and also in the end-user's network. This is to show that the Collector function could be located at all those places in the network but of course this is not necessary (e.g. not every end-user will have an instance of the Collector function). Another option, which is preferred by some of the NSPs in the ETICS consortium, is to have the Collector function located inside a trusted third party.

As an additional privacy measure, the following algorithm shall be applied when verifying that an SLA has not been violated: Whenever an SLA shall is verified for a given time period (e.g. the last hour), the end-to-

end (or edge-to-edge) metrics are evaluated and verified. If no violation is detected, this result is reported and the measurement is finished. If, on the contrary, a violation is detected, then – and only then – per-domain performance metrics are retrieved and domain(s) which have violated the SLA are identified. This algorithm also supports the scenario of deploying a trusted third party which performs the measurements.

It is expected that the transfer of monitoring data on reference points M1, M1', and M5 (cf. FIGURE 33) will consume a considerable amount of bandwidth. In order to facilitate this, an optimized, binary protocol will be proposed, the definition of which is currently work in progress. Existing protocols, such as IPFIX [RFC5101], have been considered, but they usually result in too much overhead, translating to a much higher implementation effort. Furthermore, the protocol can be redefined at a later point in time, if necessary.

It is assumed that some network operators will be quite reluctant to let additional traffic be injected into their networks only for measurement purposes. The passive network monitoring methodology accommodates this attitude perfectly. Consequently, it is assumed that the passive monitoring solution could be adopted at much more ease than solutions which foresee that traffic is injected into the network. Therefore, passive monitoring well supports fast practical deployments of the ETICS network monitoring solution.

Another substantial advantage of passive monitoring is that the actual user's traffic is measured which factually forms a "proof" of the service quality which has been delivered, while in the case of active monitoring newly generated packets are used, which might not experience the same treatment by network equipment as the users' traffic. Furthermore, with passive monitoring, it is a lot harder for a network operator to "tune" the performance evaluation results[21], because packets of the unmodified user traffic are picked for evaluation, so it not possible to provide "measurement packets" with preferential treatment.

On the negative side, it must be mentioned that very high-performance measurement equipment must be used, both in terms of CPU power and in terms of storage space directly at the probe or at least at the same site.

#### 4.1.4.4. Interfacing network monitoring with the Core ETICS system

In contrast to the core ETICS system (i.e. the part of the architecture which is not directly concerned with measurements), which can be regarded as a set of closely interrelated functions serving a common architectural purpose, the ETICS monitoring system is designed as a flexible building block that is only loosely connected to the rest of the system. The reason for such a design is threefold: Firstly, monitoring systems should in general be as independent as possible from the core system which is under measurement, in order avoid the distortion of the obtained results. Secondly, NSPs may prefer their individual selection and deployment of monitoring systems, which are not imposed by the ETICS framework. Thirdly, keeping the monitoring subsystem highly flexible and generic guarantees the highest flexibility and adaptability for all potential future measurement tasks.

Measurements can be used and integrated from various sources, e.g. OAM monitoring (cf. Section 4.1.4.3.1) and NMON (cf. Section 4.1.4.3.4).

[21] Even though we assume that a certain level of trust is established between members of the ETICS consortium.

The “SLA Monitoring” instance gets the information about contracted SLAs from the “SLA Controller”. Whenever the “SLA Monitoring” instance detects a new contracted SLA, it automatically triggers the monitoring process. This process is subject to implementation and configuration specifics, which are out of the scope of this document. All information that is necessary to perform the measurements has to be provided through the SLA. Most notably, the ingress- and egress points of the traffic as well as the description of all Points of Interconnect (PoI) in conjunction with the Interconnect Interfaces (ICI) and Traffic Delivery Points (TDP) on the path comprising the mechanisms and traffic identification methods used at the respective TDPs are required.

OAM monitoring, as already explained, is self-contained within one operator’s network. In this case, an instance of a “Continuous Monitoring Process” (CMP) functional entity is connected to the “SLA Monitoring” instance. The CMP instance coordinates retrieval of OAM monitoring information from the networking equipment.

In the case of NMON, instances of the Collector function are connected to the “SLA Monitoring” instances. In contrast to OAM monitoring, NMON spans multiple operators, ideally from edge to edge of the traffic path. To this end, additional (ETICS ASQ) routing information for all paths of all monitored SLAs will be needed. It is envisioned that this practically spans the whole Internet. Routing information within the core ETICS system is stored within the “IC Routing Controller” which is connected to “SLA Monitoring”. The latter one must be able to pass this information on to NMON and optionally to other measurement modules requiring this information.

If a contract violation is detected by the help of OAM or NMON monitoring, all pertinent information (e.g. faulty domain, violated QoS parameters, etc.) are provided to the “SLA controller” in order for the latter to react and take appropriate decision to correct the problem.

Finally, measurement results could be used by NSP to adjust their SLA offers or Network Capabilities, in particular when measurements served to fulfil traffic matrix.

### 4.1.5. The Service Enhancement Function and The Service Enhancement Functional Area

The main objectives of ETICS focus on an E2E QoS service instantiation over different service provider domains on the aggregate level and hence on managing inter-carrier ASQ paths. In addition, strategies for E2E ASQ for end-user connectivity session and service session handling are introduced in section 5.2 “End User Session Handling Architecture and Network Service API” of the ETICS deliverable D4.2 [ETICS-D4.2]. The aim of section 5.2 in D4.2 was to show how ETICS is positioned in such a wider and holistic context considering end-user ASQ connectivity session handling triggered over the “vertical” *E7* reference point and the subsequent connectivity session handling over the “horizontal” inter-NSP reference points (*E1*, *E2*, and *E3*). Moreover, the descriptions in D4.2 included how these capabilities are related to and interact with the management of inter-carrier ASQ paths. As a result, an extension to the ETICS reference architecture was presented that addresses inter-NSP session handling and Service API based interaction with Information Service Providers (InfSP).

Based on the ongoing architectural discussions within the ETICS project, the conclusion has been drawn that a more generic extension of the ETICS architecture would be beneficial. This extension, called Service Enhancement Functional Area (SEFA), enriches the basic ETICS architecture and provides the base for

individual as well as specific value added functions and services on top of the ETICS goods. Hence, the SEFA architecture can be used to enrich a broad range of ETICS services that are indicated in Section 3.2. Promising areas for extending or complementing the basic ETICS network services by means of SEFA added value services are e.g.:

- ETICS enabling end-to-end ASQ for session services (see also Section 3.3);
- ETICS enabling end-to-end ASQ for session services (see Section 3.3);
- NSP-to-Enterprise Network Services (see Section 3.2.6.2):
    - ETICS service for consumer customers,
    - ETICS services for application and content providers,
    - ETICS services for business customers.

This section puts into perspective the results of section 5.2 of D4.2 and elaborates towards a more generic ETICS architecture extension.

The Service Handling Function of Sections 5.1 and 5.2 in D4.2 has to be understood as a specific representation or implementation of the newly introduced Service Enhancement Function (SEF) of the present deliverable. In the meantime, the discussion of the extended ETICS architecture has evolved and is now able to cover and describe a broader range of added value services.

*Note: Taking this into account, the Session Handling Function (SHF) examples of D4.2, Sections 5.1 and 5.2, is assumed to be still valid with the additional constraint that the SHF is understood as a special Service Enhancement Function (SEF) within the SEFA of the extended ETICS architecture framework in D4.3.*

#### 4.1.5.1. General introduction to the SEFA concept

This section provides an introduction of the Service Enhancement Function (SEF) and Service Enhancement Functional Area (SEFA) as extension and added value enabler of the ETICS architecture. The SEF use cases Graceful Denial of Service (GDoS) and Session Handling Function (SHF) are described as examples to illustrate the capabilities of the extension.
The ETICS framework described in deliverable D2.2 includes multiple actor roles, such as edge network service provider (Edge NSP), transit network service provider (Transit NSP), transport network service provider (Transit NSP), information service provider (InfSP) and business customer (BC), as shown in FIGURE 34.

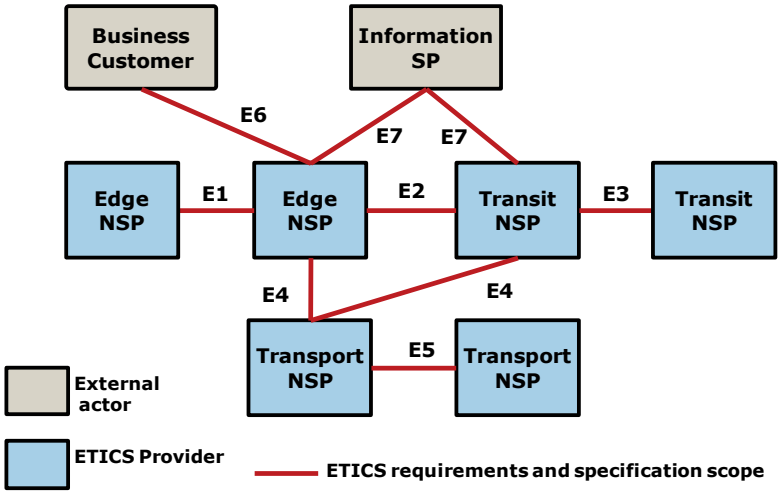


FIGURE 34: ETICS ACTOR ROLE MODEL AS PRESENTED IN D2.2

The NSP actor roles, presented in FIGURE 34, are closely related to the network plane of the ETICS architecture, described in D2.2 and D4.2. The network plane is subdivided in different sub network planes including Network Data Plane, Network Control Plane, Network Management Plane and Network Service and Business Plane. Functions within the network planes are performed by the ETICS NSPs roles for delivering and managing network or connectivity services. Besides the network plane, an application plane exists. The InfSP actor role is closely related to the application plane. The functions in the application plane are performed by the InfSP role for the purpose of delivering and managing application services. However, the application plane is not in the focus of ETICS.

The main focus of ETICS lies on the network plane. More specifically, the ETICS framework is related to establishing aggregate level resource ASQ paths for inter-carrier E2E QoS services. However, besides that ETICS has also identified additional business opportunities by enriching the basic ETICS goods and services in the application plane by means of so-called service enhancement functions (SEFs) which can be used to realize added value services that go far beyond the basic ETICS ASQ goods. The intension of ETICS is not to investigate or develop the added value aspects in the whole dimension, rather than to consider and to enable added values by means of the ETICS framework. FIGURE 35 presents the location of the service enhancement functional area in relation to the application and network plane in the ETICS architecture.

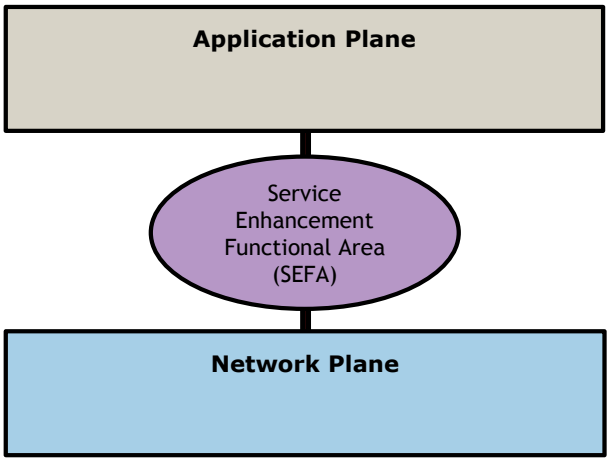


FIGURE 35: LOCATION OF SERVICE ENHANCEMENT FUNCTIONAL AREA IN RELATION TO APPLICATION AND NETWORK PLANE

The Service Enhancement Functional Area is an extension of the high level ETICS architecture in order to enable added value services and interacts between InfSP, customer and ETICS provider to enable added value functionalities. In more detail, the SEFA defines the place to enrich network plane and application plane functionalities by means of the service enhancement functions (SEFs), as shown in FIGURE 36.

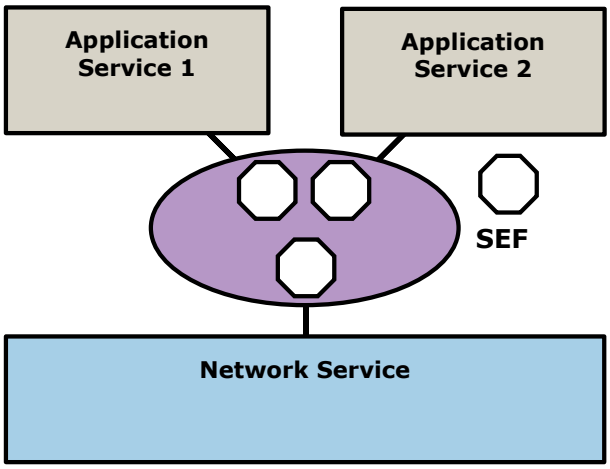


FIGURE 36: INDIVIDUAL SERVICE ENHANCEMENT FUNCTIONS AS PART OF THE SEFA

A specific added value service is represented / instantiated by means of an individual service enhancement function. This SEF gathers and combines information or parameters in order to create the specific added value service or to provide information / parameters for external added value services on top of ETICS architecture or ETICS internal services. The SEF supports both the application services and ETICS services (e.g. ASQ). The generation of added value is based on the interaction between SEF, network, application, device and customer. Moreover, each specific use case / added value service requires an individual SEF implementation and specific parameter sets.

According to the definitions of the data, network control, management planes and network service business plane in [ETICS-D2.2] and [ETICS-D4.2], a high-level definition of the Service Enhancement Functional Area (SEFA) is provided in the following:

- SEFA is an abstract area of the ETICS framework that extends the basic ETICS good (e.g. ASQ) towards specialized added value services.
- These added value services can be used to produce / enrich higher layer application services as well as network services.
- To achieve this purpose the SEFA interacts with other ETICS planes (data, control, network service business and management) in order to gather, trigger, combine and control added value service specific parameters and actions.
- The specific added value service is represented / instantiated by means of an individual Service Enhancement Function (SEF) inside the SEFA.
- SEFA is an abstract functional area that contains all possible SEFs.

In that context the Service Enhancement Function (SEF) can be understood as follows:

- The SEF represents a specific added value service or function realized by means of an individual implementation.

- In many cases the different stakeholders of a certain SEF must have special SEF relationships / agreements (regarding this SEF) that are outside the scope of the ETICS architecture.
- Besides that, it is assumed that the SEF may have value added service specific interfaces to application services.
- In general the SEF consists of three components:
    - Logic: Service Enhancement Function
    - Interfaces:
        - Internal SEFA interface: Communication between different SEF instances within the SEFA of one or multiple actors.
        - External SEFA interface: Communication between SEF/SEFA instances and information/parameter providing external units or systems, e.g. RACS [ETSI-1] or NASS [ETSI-2].

FIGURE 37 presents the ETICS reference model and reference points of the extended ETICS reference architecture. The SEFA oriented reference points or "interfaces"[22] are sketched as parallel dotted lines to the well-known ASQ related ETICS interfaces *E1* till *E7*. The name of the SEFA interfaces are equal to the ASQ related ETICS interfaces *E1* till *E7*. However, the SEFA interfaces comprises the (') as interface name extension, e.g. *E7'*. Moreover, a letter, such as *a*, *b* or *c* is used to classify different types of the "vertical" SEFA interface, as shown in FIGURE 38. This approach has been chosen in order to underline that the SEFA is an extension to the basic ETICS architecture. The interface *Ex'* is a placeholder for a generalized SEFA related interface that could be used to exchange SEF (added value) specific/related parameter sets or information. An example instantiation of an *Ex'* interface could be based on the SOAP [W3C11] protocol.

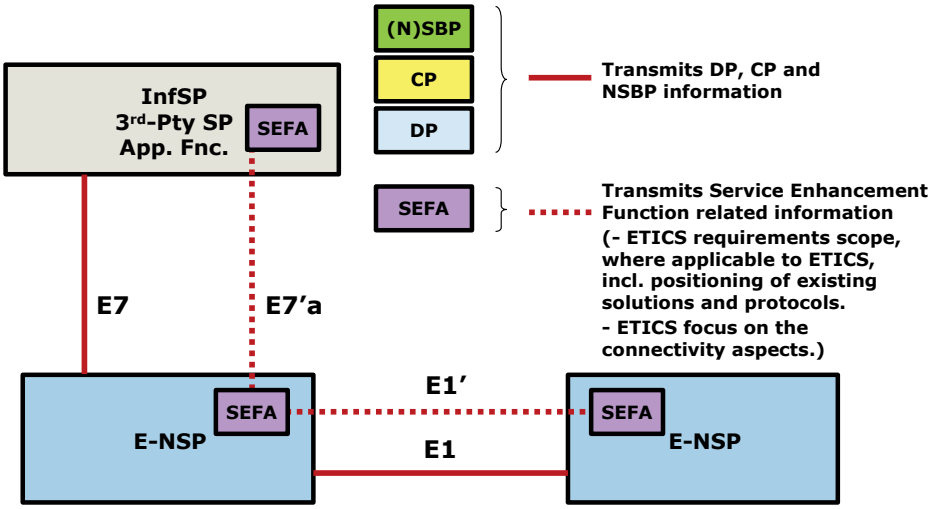


FIGURE 37: ETICS REFRENCE MODEL AND REFERENCE POINTS OF EXTENDED ETICS ARCHITECTURE

FIGURE 38 illustrates a more specific example of a value added service, a special Telco application service. Within this case, an Edge Telco Service Provider (Edge Telco SP) acts as NSP as well as value added Telco Application Service Provider and provides one or both of the two SEFA interfaces, such as *E7'a* and *E7'b* in the direction to an InfSP (3rd party). The Telco internal SEFA communication is represented by E7'c. On the

[22] The notion of reference points and interface is here used interchangeably. However, when speaking of a specific implementation of a reference point one or more technical communication interfaces are applicable. When referring to such specific communication interfaces, based on specific technical protocols, care is taken to clarify this use of the notion of interface.

other hand, the edge Telco SP has the capability to communicate with other edge Telco SP (e.g. exchange of parameter sets and information) via the "horizontal" interface *E1'* in order to realize the added value SEFA functionality for the 3rd party SP. An example added value service of such a scenario could be for instance an IMS [IMS] based service.

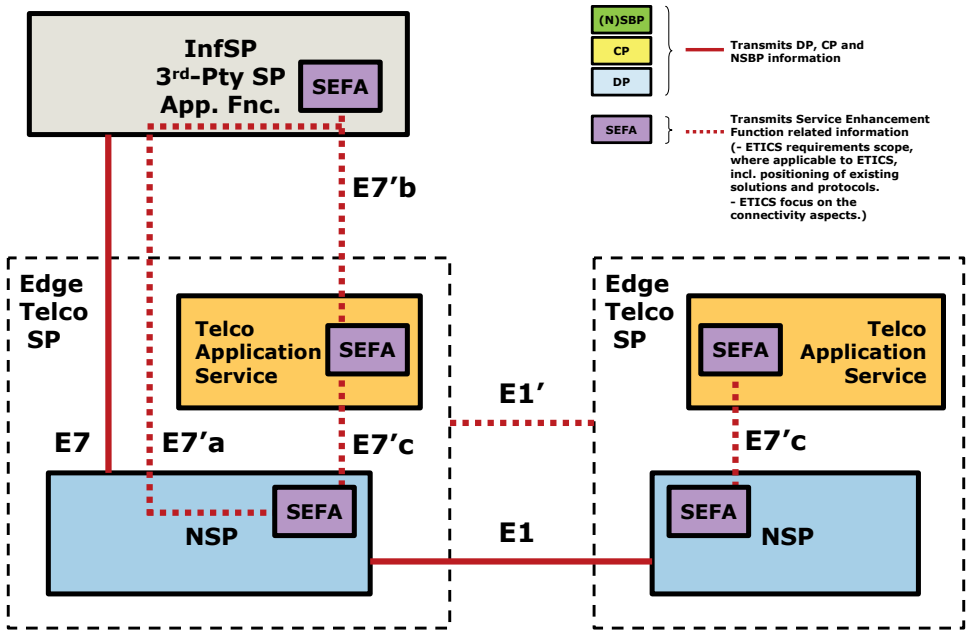


FIGURE 38: ETICS REFRENCE MODEL AND REFERENCE POINTS OF EXTENDED ETICS ARCHITECTURE - TELCO APPLICATION SERVICE PROVIDER AS EXAMPLE

Annotation to FIGURE 38:

- The SEFA communication between the two Edge Telco Service Providers (ETSP) is represented by a generalized *E1'* interface. This generalized *E1'* can be used to realize SEFA communication between either the SEFA elements of the two Telco Application Services (TAS), between the SEFA elements of the two Edge NSPs NSP, or between the SEFA elements of the TAS and the NSP. The concrete instantiation of this *E1'* depends on the SEFA functionality and how it is implemented by the Edge Telco SPs.

#### 4.1.5.2. SEFA in comparison with ITU-T NGN Reference Architecture

The Service Enhancement Functional Area (SEFA) is an abstraction across the ITU-T NGN Reference architecture [ITU2012] presented in FIGURE 39. The abstraction is across the following layers:

- Application / Service Support / Enablers
- Service Control
- Transport Control

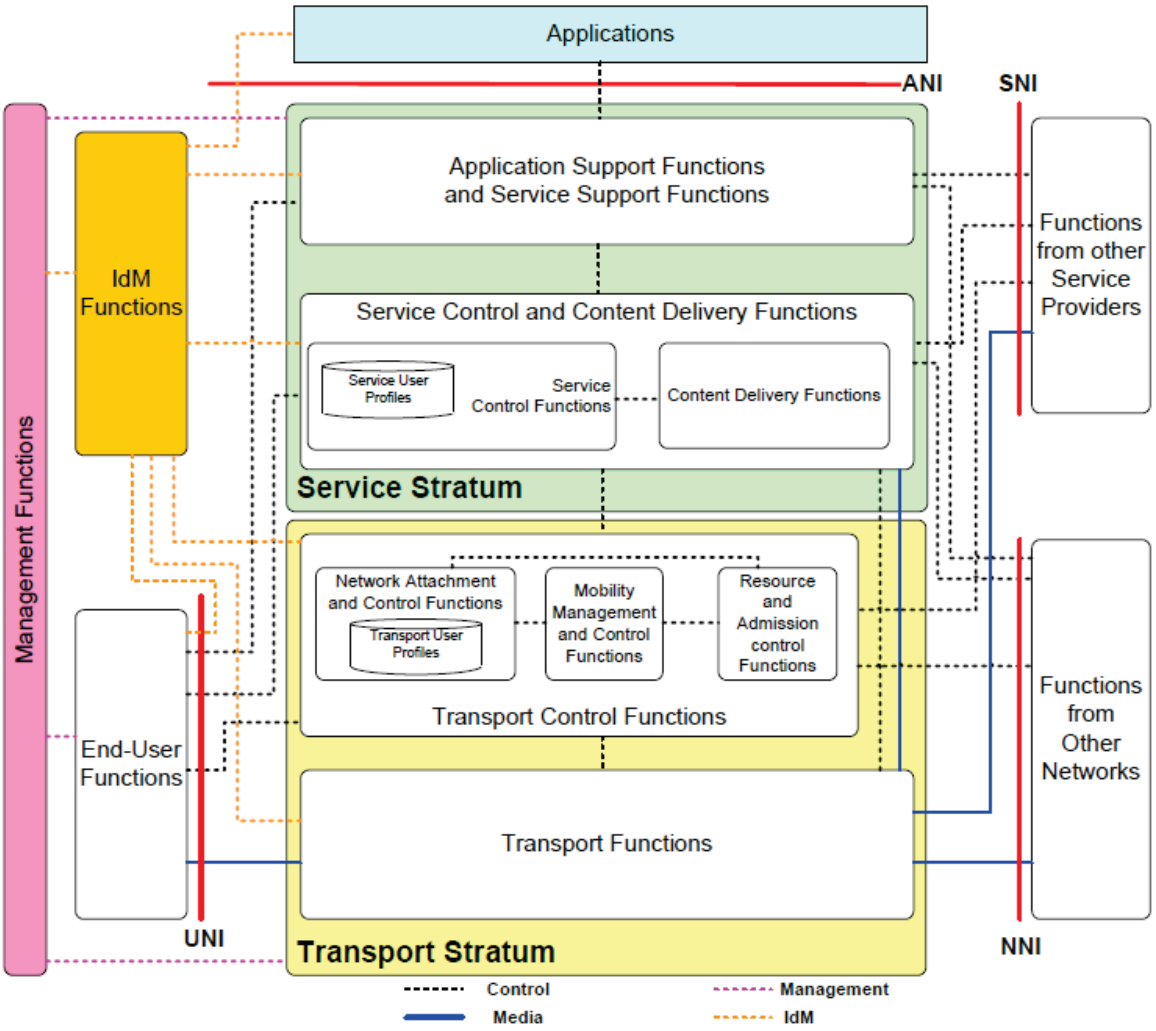


NOTE 1 – The user network interface (UNI), the network network interface (NNI), the application network interface (ANI) and the service network interface (SNI) are to be understood as general NGN reference points that can be mapped to specific physical interfaces depending on the particular physical implementations.

FIGURE 39: ITU-T REC. Y2012 - FUNCTIONAL REQUIREMENTS AND ARCHITECTURE OF NEXT GENERATION [ITU2012]

The decision for defining an ETICS specific SEFA abstraction and not simply re-using the ITU-T NGN reference architecture for the generation of added value services on top of the ETICS goods can be justified by the following statements:

- The ITU-T reference architecture is a very complex framework with lots of well-defined interfaces and functional specifications that can be to “heavy” and “over-specified” for an application in the context of the ETICS architecture.
- Depending on context the basic idea is to combine these SEFA functions in a more flexible way, and avoid strict and inflexible layering and dependencies of functions.
- Various implementation protocols can be applicable and may cover more than just one function.
- A lightweight approach such as the SEFA extension to the basic ETICS architecture may better allow for reuse of platform capabilities and reduce cost.

#### 4.1.5.3. SEFA as part of the extended ETICS architecture

The goal of individual SEF implementations is to enrich the ETICS architecture with added value functions and added value services. A possible approach to describe a right place for the individual SEFs in the ETICS architecture is to add the Service Enhancement Functional Area within the ETICS architecture. FIGURE 40 presents one possible approach to integrate the SEFA into the ETICS architecture.

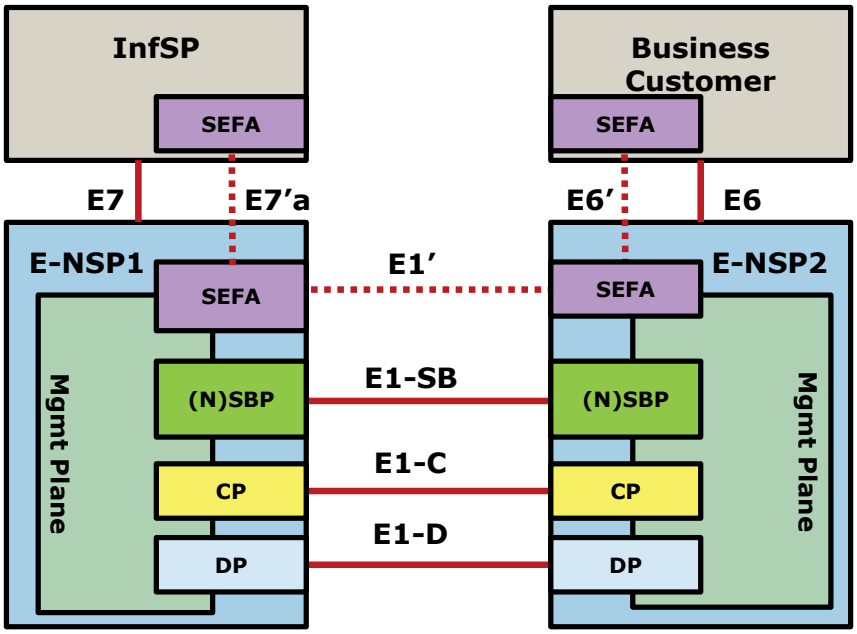


FIGURE 40: EXTENDED HIGH LEVEL ETICS ARCHITECTURE INCLUDING SERVICE ENHANCEMENT FUNCTIONAL AREA (SEFA) FOR ADDED VALUE SERVICES

However, no final decision has been made regarding the more detailed positioning of the SEFA and specific SEFs in relation to the ETICS network planes and this still has to be considered as work in progress.

##### *4.1.5.3.1. SEFA and related Technical Requirements of D2.2*

As described in Section 4.1.5.1, the SEFA interacts with InfSP, customer and NSP and is mainly active on the vertical interfaces, such as *E7'* and *E6'* (in certain scenarios / added value services the SEF may use as well horizontal interfaces, such as *E1'*). That is why the realisation of value added services based on SEFs depends on the compliance of the ETICS architecture with the requirements for vertical interfaces as described in the technical requirements of [ETICS-D2.2]. Especially the following requirements from [ETICS-D2.2] have to be fulfilled:

| REQ NBR | Description | Relation to SEFA / SEF |
|---|---|---|
| BR-VERT-02 | The ASQIEC shall allow the Buyer to check on-demand if the end-user at the session end point has required privileges to receive ASQ Interactive end-user connectivity | Example SEF use case – Graceful Denial of Serivce (GDoS). BR-VERT-02 can be used to verify that the customer is allowed to use high quality data transport. |

| | | |
|---|---|---|
| BR-VERT-03 | The ASQIEC shall allow charging based on different quality requirements and volumes in the two traffic directions. | Example SEF use case – Graceful Denial of Serivce (GDoS). BR-VERT-03 can be used to provide differentiated charging. |
| BR-VERT-08 | The ASQIEC shall allow the NSP to coordinate the inter-NSP resources and the admission control needed in order to include end-points in domains beyond the neighbour NSP. | Example SEF use case – Graceful Denial of Service (GDoS). BR-VERT-08 can be used to request network capabilities of the customer connecting NSP. |
| BR-VERT-10 | ASQECC shall allow the Buyer to check on-demand if the user at the session end point has the required connectivity service privileges to receive ASQ content delivery from the Buyer. | Example SEF use case – Graceful Denial of Service (GDoS). BR-VERT-10 can be used to verify that the customer is allowed to use the SEF added value service. |
| BR-VERT-11 | ASQECC shall allow the Buyer to update the end-user access connectivity profile controlled by the edge NSP per end-user provisioning or session invocation time scales, according to policies of the edge NSP. | Example SEF use case – Graceful Denial of Service (GDoS). BR-VERT-11 can be used to request the parameters of the current customer access network capabilities from the edge NSP. |
| TR-SLA-GEN-10 | The SLA framework must/should include at least two levels, pertaining to sessions and aggregated flows respectively. | This requirement can be used to generate SLA for added value services on top of SLA dedicated to data transport services session (e.g. ASQ) |
| TR-SLA-GEN-11 | The SLA framework must provide a mechanism to set-up and tear-down a session for the purpose of providing a service to the user. | This requirement forms the basis for on demand value added services on top of data transport service session. |
| TR-SLA-GEN-12 | The SLA framework must provide a mechanism to modify the parameters of an active session | This requirement can be used by a specific SEF based added value service in order to adjust the session parameters of an aggregated session in an underlying ASQ to meet the application service requirements. |
| TR-NET-AC-02 | The admission control process must be able to perform admission control for each resource / session / SLA request independently. | This requirement can be used by a specific SEF in order to provide admission controlled added value services on top of an SLA based ETICS good. |
| TR-NET-AC-04 | The choice for the specific (intra-provider) admission control process and for the needed and implemented | The SEF has to provide a NSP admission control independent interface in order to create admission controlled added value |

| | | |
|---|---|---|
| | mechanisms (which can work on different abstraction layers) has to be left to the NSP. There exists no "One size fits all" mechanism. | services. |
| TR-NET-AC-06 | It must be possible for a service provider to trigger the admission control process during different phases of the SLA / service lifecycle process. However, within the context of the ETICS project, admission control is assumed to be mainly performed during the "SLA provisioning / invocation" phase of the SLA / service lifecycle process. | The SEF will use this requirement in order trigger or adjust added value services during the runtime of existing ETICS SLA lifecycle. |

#### 4.1.5.4. Examples for Service Enhancement Function use cases – GDoS and SHF

After this very theoretical discussion about the structure of the SEFA and how the SEFA can be represented within an extended ETICS architecture framework, the following subsection is devoted to the description and discussion of two exemplary real SEF use cases – "Graceful Denial of Service" (GDoS) and "Session Handling Function" (SHF).

Note: At this point it should be once more mentioned that the ongoing discussions regarding the ETICS architecture have updated the understanding of the Session Handling Functions as introduced in D4.2. From a current point of view, the SHF can be understood as special Service Enhancement Function that is dedicated to add a service-related and / or connectivity-related value to an existing ASQ session. In that terminology, D4.2 has created the terms SHF-S and SHF-C and illustrated with three cases how either a SHF-S respectively SHF-S and SHF-C can be implemented and used by different actors in specific network scenarios (Network Service Provider, Combined Network Service and Communication Service Provider, Home Network Interaction).

The present section picks up again the first two cases of [ETICS-D4.2] and highlights how these cases fit into the SEFA architecture.

##### *4.1.5.4.1. Example SEF use case - Session Handling Function (SEF-SHF)*

In the following, the use case "Service Enhancement Function – Session Handling Function" (SEF-SHF) is described. This use case requires information from the network provider to gather all parameters to realise the specific SEF-SHF added value service. In D4.2, the use case SEF-SHF has been described by means of the terminology Session Handling Function (SHF). However, the further development of the ETICS architecture towards a more generic view on added value services by means of the introduced SEFA needs a reworking of the description of use case 1 'network service provider role only' in Section 5.2.2.1 of D4.2. In the following the reworked architecture of the SEF-SHF by means of the SEFA terminology is presented.

**CASE 1: Network Service Provider role only**

In the following scenario the network session handling for enabling an end-to-end ASQ for end-user connectivity is described. FIGURE 41 presents the reference architecture of the network service provider role within this scenario. The relationship between NSP and InfSP is realised by means of interfaces *E7* and *E7'* and is possible in two ways:

- InfSP is owned by the NSP (network operator) itself. InfSP processes are assumed to be harmonised with NSP process.
- InfSP is separated from the NSP and owned by a different actor. The InfSP processes are not harmonised with the NSP processes.

In order to deal with both variants, the reference architecture envisions specific functional elements, such as Connectivity Profile Function (CPF), Traffic Engineering Function (TEF), Resource and Admission Control Function (RACF) and Service Enhancement Function - Session Handling Function (SEF-SHF) to perform end-to-end QoS admission control and to provide end-to-end QoS handling for end-user connectivity sessions. The following functions are the key elements of the reference architecture:

- Connectivity Profile Function (CPF):

  The Connectivity Profile Function (CPF) stores information about users' network access and connectivity capabilities, such as for example maximum granted bandwidth, QoS parameters, limitations of maximum login time and volume of traffic. Moreover, device capabilities, such as available network interfaces (e.g. WLAN, Ethernet, Bluetooth, etc.), processor power, maximum display resolution, type of electric power supply (battery or AC power supply unit) could be stored in the CPF and be used for end-user management as well.

  Besides that, the CPF stores also accounting information and provides the base for billing, statistics and policy management.

  The CPF will be updated in a predefined cycle.

  The CPF communicates with the Service Enhancement Function - Session Handling Function (SEF-SHF).

- Traffic Engineering Function (TEF):

  The on-line Traffic Engineering Function (TEF) is a crucial function and is expected to bridge the aggregate level resources with the end-user connectivity level by monitoring the network resource situation and managing including updating traffic steering policies and the ASQ paths, both the inter-NSP and the intra-NSP ones. It may also perform path computation for end-to-end ASQ for end-user connectivity sessions that have high demands, while the typical case is that the routing of end-user connectivity flows are based on traffic steering policies. By some means the SEF-SHF is directed by the traffic steering policies that the TEF manages.

  The TEF communicates with the Service Enhancement Function - Session Handling Function (SEF-SHF).

- Resource and Admission Control Function (RACF):

The Resource and Admission Control Function (RACF) provides policy based transport control that allows to request and reserve transport resources from the transport networks. The RACF ensures that requested network resources are available in the corresponding access (transport) network and that the customer is allowed to use these resources.

The RACF communicates with the Service Enhancement Function - Session Handling Function (SEF-SHF).

The RACF could be based on the Resource and Admission Control Subsystem (RACS) framework specified by ETSI TISPAN in ETSI ES 282 003 [ETSI-1].

- Service Enhancement Function - Session Handling Function (SEF-SHF):

  The Service Enhancement Function - Session Handling Function (SEF-SHF) is part of the Service Enhancement Functional Area (SEFA) and acts as a specific implementation of the generic described Service Enhancement Function (SEF) between network plane and CPF to enable mutual awareness of network connectivity and user profile issues. The SEF-SHF as mediator translates and determines required information of the requested service. This means the SEF-SHF investigates selected information obtained from the network and CPF to derive input parameters or control information to setup ETICS services, such as end-to-end ASQ end-user connectivity. The derived input parameters can be transmitted via the interfaces *E1'*, *E2'* as well as *E7'*.

  The SEF-SHF is a functional element that has to be designed depending on the required functionality needed to perform or to setup an ETICS services, such as end-to-end ASQ of an end-user. The implementation depends on the focused functionality needed.

  The SEF-SHF can be applied to perform investigations on demand in the case of connectivity session initiation process. Moreover, the SEF-SHF can perform investigation during an active connectivity session. The CPF is updated during every SEF-SHF process cycle.

  The SEF-SHF communicates provider network internal with the CPF, TEF and the RACFs. Network external the SEF-SHF communicates with other edge NSPs via interface *E1'* and with transit NSPs via interface *E2'*. The interface *E7'* is used to exchange information with an information service provider. The interaction between SEF-SHF and network internal functional elements as well as network external functional elements of other operators takes place according to the designed workflow.

  With regard to the inter-carrier ASQ path establishment, the SEF-SHF has to be compatible and thus aware of the inter-carrier ASQ establishment process. As a result, by means of the SEF-SHF the QoS and ASQ capabilities can NSP-internally be abstracted, mapped or simplified to be exchanged via the *E1'*, *E2'* or *E7'* interface. Hence, the SEF-SHF as part of the SEFA (and acts as mediator) can build up the *E1'*, *E2'* or *E7'* interface depending on the required level of inter-NSP and InfSPs information exchange, however, keeping in mind the ETICS architecture conformity.

The functional elements, such as Connectivity Profile Function, Traffic Engineering Function, Resource and Admission Control Function and Service Enhancement Function - Session Handling Function are independent from each other, but all these elements can communicate independently with each other.

Note that the Transit NSP is optional, depending on the specific case in focus.

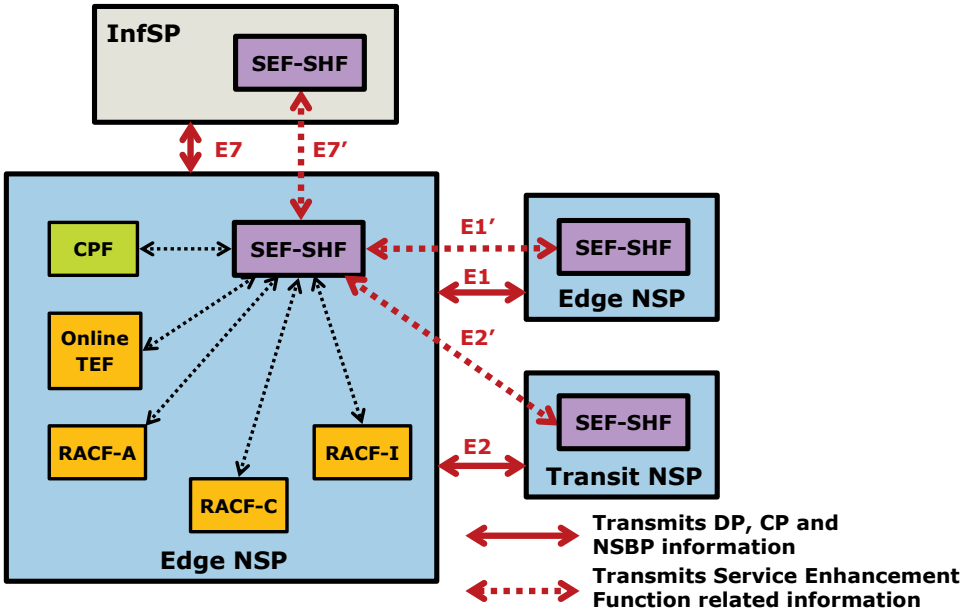


FIGURE 41: NSP ROLE ONLY

As shown in FIGURE 41, the NSPs and also their SEF-SHFs interact by means of the interface *E1'* with other edge NSPs and via the interface *E2'* with transit NSPs and their SEF-SHFs. *E1'* is a parallel interface to interface *E1* containing connectivity related information and parameters. The interface *E1'* could contain information, e.g. user policies, ASQ information, point of attachment, point of delivery and derived control information from the SEF-SHF.

The interface *E7'* is used to exchange information between the network plane and the application plane in order to derive added value control information, as shown in FIGURE 41 above. Transmitted information can be, e.g. application requirements, network parameters or derived control information from the SEF-SHF.

The design of interface *E1'*, *E2'* and *E7'* has to be extensible and must not be limited to a specific amount of parameters.

A candidate protocol for the interfaces *E1* and *E2* could be the Diameter protocol as described in [ETSI-1] as interface Rr. Candidate protocols or frameworks for the *E1'*, *E2'* and *E7'* could be, e.g. Simple Object Access Protocol (SOAP) [W3C11].

#### *4.1.5.4.2. Example SEF use case – Graceful Denial of Service (SEF-GDoS)*

In the following, the use case "Service Enhancement Function – Graceful Denial of Service" (SEF-GDoS) is described. This use case requires information from the network provider as well as from the communication / information service provider to gather all parameters to realise the specific SEF-GDoS added value service. In D4.2, the use case SEF-GDoS has been described by means of the terminology Session Handling Function (SHF). However, the further development of the ETICS architecture towards a more generic view on added value services by means of the introduced SEFA needs a reworking of the description of use case 2 'combined network service provider and communication service provider role' in

Section 5.2.2.2 of D4.2. In the following the reworked architecture of the SEF-GDoS by means of the SEFA terminology is presented.

**CASE 2: COMBINED NETWORK SERVICE PROVIDER AND COMMUNICATION SERVICE PROVIDER ROLE**

FIGURE 42 presents the reference architecture of the combined network service provider and communication service provider role scenario. The reference architecture is based on the functional elements described in FIGURE 41. Additional functions compared to the network service provider only role scenario are Service Profile Function (SPF) and Service Enhancement Function – Graceful Denial of Service (SEF-GDoS) to perform application service session handling besides e2e QoS admission control for end-user connectivity. The following additional functions extend the number of key elements of the reference architecture described in Section 4.1.1. As noted above, many use cases applies when considering application level aspects; the one below is just an example and the details of the potential interfaces will (may) only be studied further for some use cases. However, such analysis is beyond this deliverable.

- Service Profile Function (SPF):

  The Service Profile Function (SPF) stores information about subscribed services of the user. It is possible to store service related parameters or information, such as critical level of available network resources for service provisioning in user satisfying quality, in the SPF. These parameters could be used by the Service Enhancement Function (SEF) to derive input parameters for the application server with the aim, e.g., to adapt service quality (e.g. resolution of video) that is provideable by the network.

  The SPF will be updated every time after an SEF process cycle.

  The SPF communicates with the Service Enhancement Function.

- Service Enhancement Function – Graceful Denial of Service (SEF-GDoS):

  The Service Enhancement Function – Graceful Denial of Service (SEF-GDoS) is an added value function between application plane, SPF and network plane that enables mutual awareness of each other. The SEF-GDoS collects, evaluates and utilises service related information, e.g. user service policy, service requirements and network performance to derive input parameters or control information for the application service session initiation, e.g. service is provideable in requested service quality. Within the SEF-GDoS the service related part of the SEF (SEF-S) (formerly in D4.2 called SHF-S) interacts with the connectivity related part of the SEF (SEF-C) (formerly in D4.2 called SHF-C) in order to retrieve network input parameters.

  The SEF-GDoS also provides / determines service related information requested from the 3rd Party Provider / Value Added Service Provider. The 3rd Party Provider communicates via an *E7'* interface with the Information Service Provider. The InfSP translates or forwards the requested information of the 3rd Party Provider to the NSP and their SEF-GDoS by means of the interface *E7'*. Vice versa the SEF-GDoS sends the application service control information via the *E7'* to the InfSP and via the *E7'* interface to the 3rd Party Provider to control the session initiation. Moreover, derived

parameters or requested parameters can be transmitted via the interfaces *E1'* and *E2'* to external network providers.

The SEF-GDoS is a functional element that has to be designed depending on the required functionality needed to initiate and control an application service session. The implementation depends on the focused functionality needed.

The SEF-GDoS can be applied to perform investigations on demand in the case of application service session initiation process. Moreover, the SEF-GDoS can perform investigation during an active application service session. The SPF is updated by every SEF-GDoS process cycle.

The service related part of SEF-GDoS communicates provider network-internally with the SPF and the connectivity related SEF-GDoS part. Externally, the SEF-GDoS exchanges information with other edge NSPs via interface *E1'* and with transit NSPs via interface *E2'*. The interaction between service related part of SEF-GDoS, SPF, 3rd Party Provider, connectivity related SEF-GDoS part and network external functional elements of other operators take place according to the designed workflow. As described in subsection 4.1.5.4.1 under the topic SEF-SHF, the SEF-GDoS has to be compatible and thus aware of the inter-carrier ASQ establishment process as well. This means that the service related part of SEF-GDoS has to be harmonised with the connectivity related part of SEF-GDoS and thus with the ETICS architecture. Due to the (direct) interaction between the service and connectivity related parts within the SEF-GDoS, as illustrated in FIGURE 42, the evaluation of application / business requirements provided by the InfSP / 3rd party provider regarding the ASQ capabilities of the edge NSP should be possible, leading to, e.g., indicators used for application service initiation. Moreover, the interaction of service and connectivity related parts of the SEF-GDoS will be useful with regard to abstraction, mapping or simplification of exchanged ASQ path setup information via the interface *E1'* and *E2'*. As a result, the SEF-GDoS provides the base for business and service enabler that would not appear without the consideration of application requirements in the ETICS architecture.

The functional elements, such as Service Profile Function, service related part of SEF-GDoS (SEF-S) and connectivity related part of SEF-GDoS (SEF-C) are independent from each other, but all these elements can independently communicate with each other.

Note that the Transit NSP is optional, depending the specific case in focus.

The NSP (and its SEF-GDoS) interact via the interface *E1'* with other edge NSPs and via the interface *E2'* with transit NSPs (and their SEF-GDoS). *E1'* is a parallel interface to the interface *E1* signalling service and connectivity related information and parameters as for instance business model related information, such as ASQ quota and user privileges.

Moreover, the edge NSP or CmSP interact via *E7* and *E7'* with Information Service Providers and via *E6* and *E6'* with Business and End Customers. In this case, the interface *E7'* is used to exchange added value information between the combined NSP/Communication SP role and the application plane, as shown in FIGURE 42, the transmitted information can be, e.g., application requirements, network parameters or derived application service session control information from the SEF-GDoS. We also recognize that the InfSP may offer services to his 3rd party providers. By means of an *E7'* interface, a 3rd Party Provider is able to request or receive application service session control information to setup or initiate an 3rd Party

application service based on users demand. In general, the NSP/CmSP should not need to know about any subsequent offerings by the InfSP, and hence, such "cascaded" direct dependencies can be avoided, however, this possibility is an enabler for new application services as well as for new business.

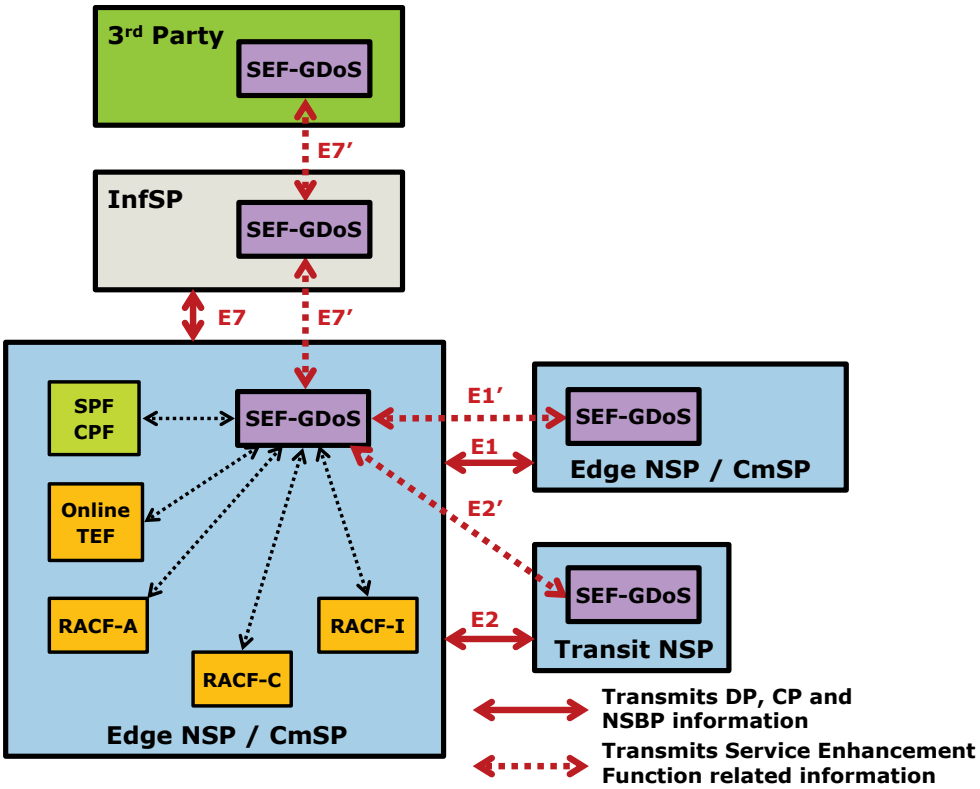


FIGURE 42: COMBINED NSP AND COMMUNICATION SP (CMSP) ROLE IN SEFA USE CASE

The design of interface *E1'*, *E2'* and *E7'* has to be extensible and must not be limited to a specific amount of parameters.

A candidate protocol for the interfaces *E1'* and *E2'* could be the Diameter protocol as described in [RACS.2010] as interface Rr. In addition to the Diameter protocol, the Session Initiation Protocol (SIP) as used in the IP Multimedia Subsystem could be a candidate protocol for *E1'* as well. Candidate protocols or frameworks for the *E7'* could be, e.g. the Simple Object Access Protocol (SOAP) [W3C11], Parlay X [PaX], GSMA OneAPI [OnAp], Open Mobile Alliance (OMA) Service APIs [OMA], Wholesale Application Community (WAC) [WAC], JAIN and Java in Communications (2004) [Or11], [JaCo11], or JAIN and Open Networks (2003) [Or11].

**General GDoS service description**

The use case "Graceful Denial of Service" (GDoS) enables an "engaged" signal for IP services considering available resources in the user access network. As a result, avoidance of reduced service quality due to over-booking the available bandwidth in the user access network is achieved. FIGURE 43 provides a high level view on the graceful denial of service use case scenario. The "engaged" signal for IP services is enabled by means of considering the available quality units in user access profile of the access network.

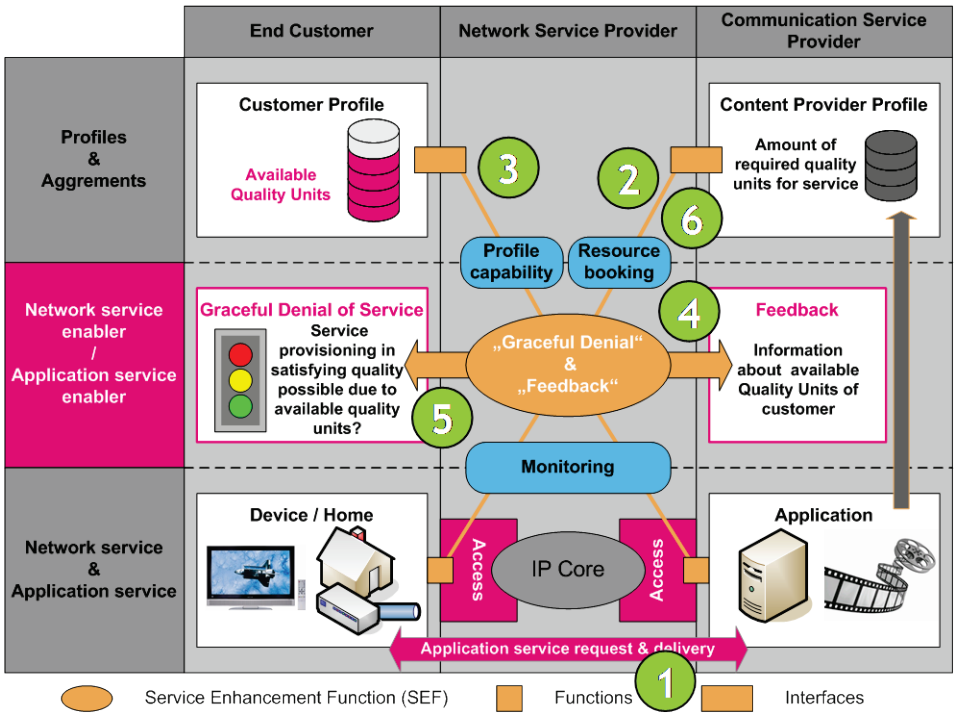


FIGURE 43: SEF USE CASE GRACEFUL DENIAL OF SERVICE

The high level description of GDoS can be given as follows:

1. Customer requests application service, such as IPTV, in high quality from communication service provider (CmSP).
2. CmSP sends required amount of quality units (needed for providing requested application quality) to the network service provider (NSP).
3. Service Enhancement Function (SEF) requests available quality units from customer profile.
4. SEF - Graceful Denial of Service evaluates capabilities of the customer profile, such as available quality units and provides feedback to the CmSP.
5. CmSP informs customer about the state of service provisioning capabilities, such as whether the requested service is provideable in requested quality or not.
6. CmSP carries out resource ordering from the NSP in the case of user confirms start of application service.

### GDoS Message Flow Diagram

The presented high-level view on the Graceful Denial of Service use case scenario in FIGURE 43 is described by means of the flow diagram in FIGURE 44 in more detail. The flow diagram depicts the different actor roles involved in the GDoS scenarios. Moreover, the sequential steps needed to perform the “engaged” signal for the IP service are presented. Service related information of the SEF-GDoS is processed by the SEF-S and connectivity related information of the SEF-GDoS is processed by the SEF-C.

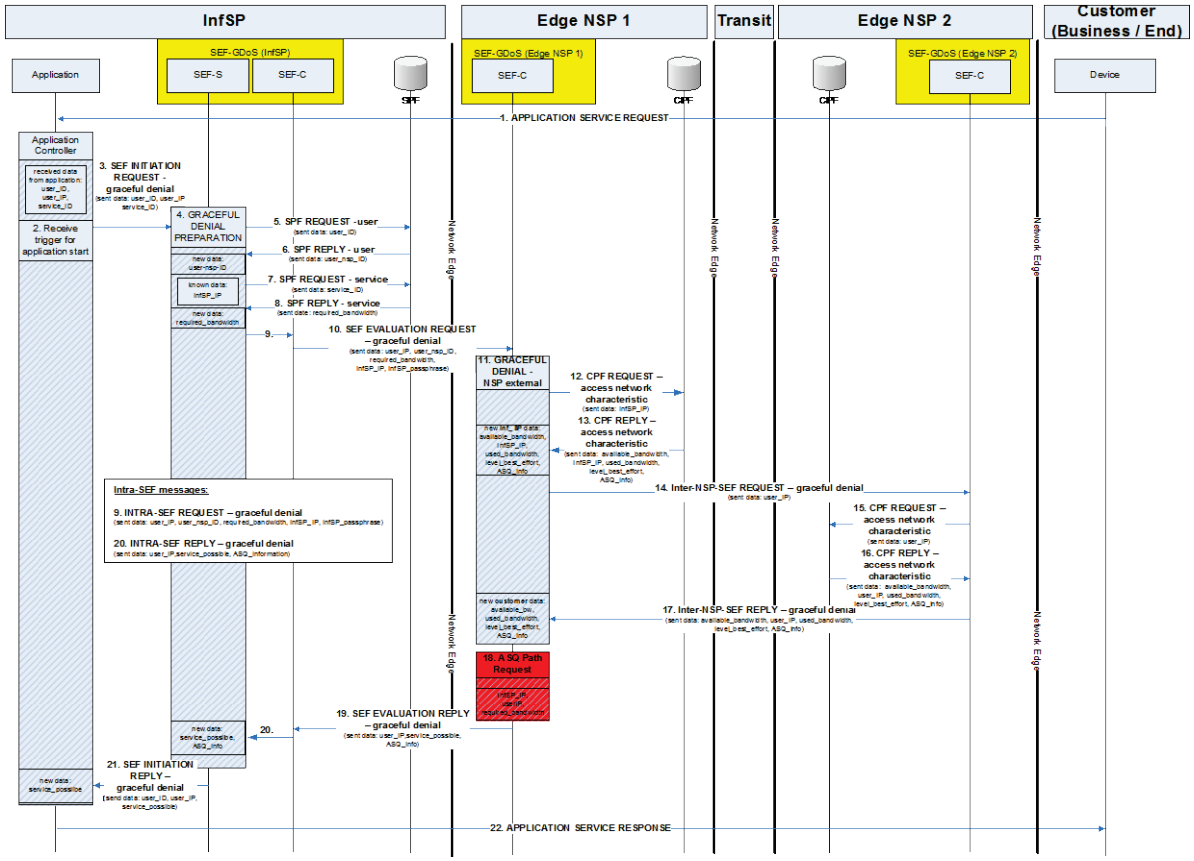


FIGURE 44: SEF-GDOS FLOW DIAGRAM

Please, note that FIGURE 44 is represented in full scale in the Annex Section 8.1.

The sequential steps of the flow diagram are presented in the following:

1. **Customer <> Information Service Provider (InfSP):** <APPLICATION SERVICE REQUEST> device requests high quality (ASQ based) application service, e.g. video service.
2. **InfSP**: <Receive trigger for application start> application controller (AC) recognizes service request of device/customer.
3. **InfSP**: <SEF INITIATION REQUEST - graceful denial> AC sends “SEF INITIATION REQUEST – graceful denial” to start the evaluation process of available resources in customer access network. (sent data: user_ID, user_IP, service_ID)
4. **InfSP**: <GRACEFUL DENIAL PREPARATION> Service Enhancement Function – Service (SEF-S) initiates “GRACEFUL DENIAL PREPARATION” with the aim to request user network access and service related information.
   (received data: user_ID, user_IP, service_ID)
5. **InfSP**: <SPF REQUEST - user > Service Enhancement Function – Service (SEF-S) requests by means of “SPF REQUEST – user” user related point of attachment information from the Service Profile Function (SPF).
   (sent data: user_ID)
6. **InfSP**: <SPF REPLY - user> SPF provides by means of “SPF REPLY – user” point of attachment information, such Edge NSP2 to SEF-S.
   (sent data: user_nsp_ID)

7. **InfSP**: <SPF REQUEST - service> SEF-S requests by means of “SPF REQUEST – service” service related information from the Service Profile Function (SPF).
   (sent data: service_ID)
8. **InfSP**: <SPF REPLY - service> SPF provides by means of “SPF REPLY – service” service requirements, such as required bandwidth to SEF-S.
   (sent data: required_bandwidth)
9. **InfSP**: <INTRA-SEF REQUEST – graceful denial> SEF-S transmits parameters needed for “GRACEFUL DENIAL – NSP internal” process by means of “INTRA-SEF REQUEST – graceful denial” to Service Enhancement Function – Connectivity (SEF-C).
   (sent data: user_IP, user_nsp_ID, required_bandwidth, InfSP_IP, InfSP_passphrase)
10. **InfSP <> NSP1**: <SEF EVALUATION REQUEST – graceful denial> SEF-C initiates by means of “SEF EVALUATION REQUEST – graceful denial” the graceful denial of service process at NSP1.
    (sent data: user_IP, user_nsp_ID, required_bandwidth, InfSP_IP, InfSP_passphrase)
11. **NSP1**: <GRACEFUL DENIAL - NSP external> SEF-C initiates “GRACEFUL DENIAL – NSP external” investigating NSP2 connected users network access capabilities.
12. **NSP1**: <CPF REQUEST – access network characteristic> SEF-C requests by means of “CPF REQUEST – access network characteristic” available network access resources of InfSP from CPF.
    (sent data: InfSP_IP)
13. **NSP1**: <CPF REPLY – access network characteristic> CPF provides by means of “CPF REPLY – access network characteristic” InfSP network access information to the SEF-C.
    (sent data: available_bandwidth, InfSP_IP, used_bandwidth, level_best_effort, ASQ_info)
14. **NSP1 <> NSP2**: <Inter-NSP-SEF REQUEST – graceful denial> SEF-C requests by means of “Inter-NSP-SEF REQUEST – graceful denial” available network access resources of user connected to NSP2.
    (sent data: user_IP)
15. **NSP2**: <CPF REQUEST – access network characteristic> SEF-C requests by means of “CPF REQUEST – access network characteristic” available network access resources of user from CPF.
    (sent data: user_IP)
16. **NSP2**: <CPF REPLY – access network characteristic> CPF provides by means of “CPF REPLY – access network characteristics” user network access information to the SEF-C.
    (sent data: available_bandwidth, user_IP, used_bandwidth, level_best_effort, ASQ_info)
17. **NSP2 <> NSP1**: <Inter-NSP-SEF REPLY – graceful denial> SEF-C provides by means of “Inter-NSP-SEF REPLY – graceful denial” available network access resources of NSP2 connected user to “GRACEFUL DENIAL - NSP external” at NSP1.
    (sent data: available_bandwidth, user_IP, used_bandwidth, level_best_effort, ASQ_info)
18. **NSP1**: <ASQ Path Request> In the case of non established ASQ path (ASQ_info=0) from InfSP to customer the “ASQ path request” is used to initiate ASQ path establishment process in ETICS community.
    (sent data: InfSP_IP, user_IP, required_bandwidth)
19. **NSP1 <> InfSP**: <SEF EVALUATION REPLY – graceful denial> SEF-C provides results of “GRACEFUL DENIAL – NSP external” by means of “SEF EVALUATION REPLY – graceful denial” to SEF-C of InfSP if user network access resources are sufficient enough to provide service in required quality.
    (sent data: user_IP, service_possible, ASQ_info)

20. **InfSP**: <INTRA-SEF REPLY – graceful denial> SEF-C transmits results of "GRACEFUL DENIAL - NSP external" by means of "INTRA-SEF REPLY – graceful denial" to SEF-S.
    (sent data: user_IP, service_possible, ASQ_info)
21. **InfSP**: <SEF INITIATION REPLY – graceful denial> SEF-S transfers the "GRACEFUL DENIAL – NSP external" result (if user network access resources are sufficient enough to provide service in required quality) by means of "SEF INITIATION REPLY – graceful denial" to the application controller.
    (sent data: user_ID, user_IP, service_possible)
22. **InfSP**: <APPLICATION SERVICE RESPONSE> Application controller uses "GRACEFUL DENIAL – NSP external" result to start service and/or to provide feedback to the customer.

## 4.2. ETICS FEATURES OR DEPLOYMENT SCENARIOS

Deliverable D4.2 has identified six different scenarios for the ETICS service composition. Although each of these scenarios implies a different communication pattern between NSPs in the ETICS community, they can be seen as different features of the same architecture. Each of these features fits a different context. For instance, the on-demand feature can fit very customized or uncommon ASQ requests; it is therefore more suitable for a starting market. On the other hand, the pre-computed scenario is more suitable for a mature market where NSPs know what kind of offers can best fit the ETICS customers' demands. A prioritisation of the different options was performed from the business perspective (within [ETICS-D3.3]) and will be presented in Section 5.3.

In this section, we present a high level specification for each of these features. Each section deals with one scenario by applying the following three steps: first, it recalls the corresponding ETICS architecture from deliverable D4.2. On this basis, it subsequently presents a high level UML sequence diagram that captures the inter-NSP interactions. Finally, it explores ways to implement the scenarios.

### 4.2.1. ON-DEMAND (PULL) OFFER WITH UNIQUE CENTRALIZED COMPOSITION ENTITY SCENARIO

#### 4.2.1.1. Architecture High-level View

FIGURE 45 retakes the ETICS service and business plane architecture in the case of on-demand scenario, with NSP chain learning and a centralized entity that handles the service composition process for an ETICS community of NSPs. This centralized entity works as a facilitator and acts as the interface that will receive all the ETICS customers' requests. Upon the reception of a request for an inter-carrier ASQ path, the central entity performs the service composition by requesting (*pulling*) offers from the set of NSPs that are likely to satisfy the customer request. Once done, the central entity composes these offers and creates a single offer that it proposes to the ETICS customer. The latter might order it if it finds it satisfying.

The NSPs in the community define and agree on the mission of this facilitator and they fix its objective. This objective can be technical, business, or both. The goal of this section, and of the subsequent design of the architecture, is to mainly focus on the technical possibilities and leave the business considerations as open as possible. The facilitator can be either a neutral entity that is used as a technical tool that works for the benefit of the overall community; it can also be a third party that works for its own benefit. The business analysis of the different features will be separately analysed in Section 5.3.

In this first feature (scenario), the service and business planes contain mainly four entities. We first detail their roles and interactions.

- **Service discovery facilitator**: This entity that may interact with the control plane (in a fully automated approach) to obtain information that is needed for NSP chain learning. We have exposed in Section 3.3 levels of detailed information that could be used by this entity. The collected intra-domain information will be made available to the central facilitator entity, which will use it to perform the NSP chain learning.
- **SLA requests processor**: This entity interacts with the centralized facilitator entity, from which it receives requests to provide specific offers. Upon the reception of such a request, the SLA requests processor triggers the product offer creation that will be done by the «product offer creator» entity.
- **Product offer creator**: This entity may interact with the control plane and data plane or indirectly through OSS north interface. It makes the resource inventory and creates offers upon the reception of a request from the SLA-requests processor entity. With respect to the SLA lifecycle, the offer is specified by an SLA instance.
- **Service instance manager:** instantiates the offers upon the reception of a request. It also interacts with the billing/accounting entity to charge for the offer.

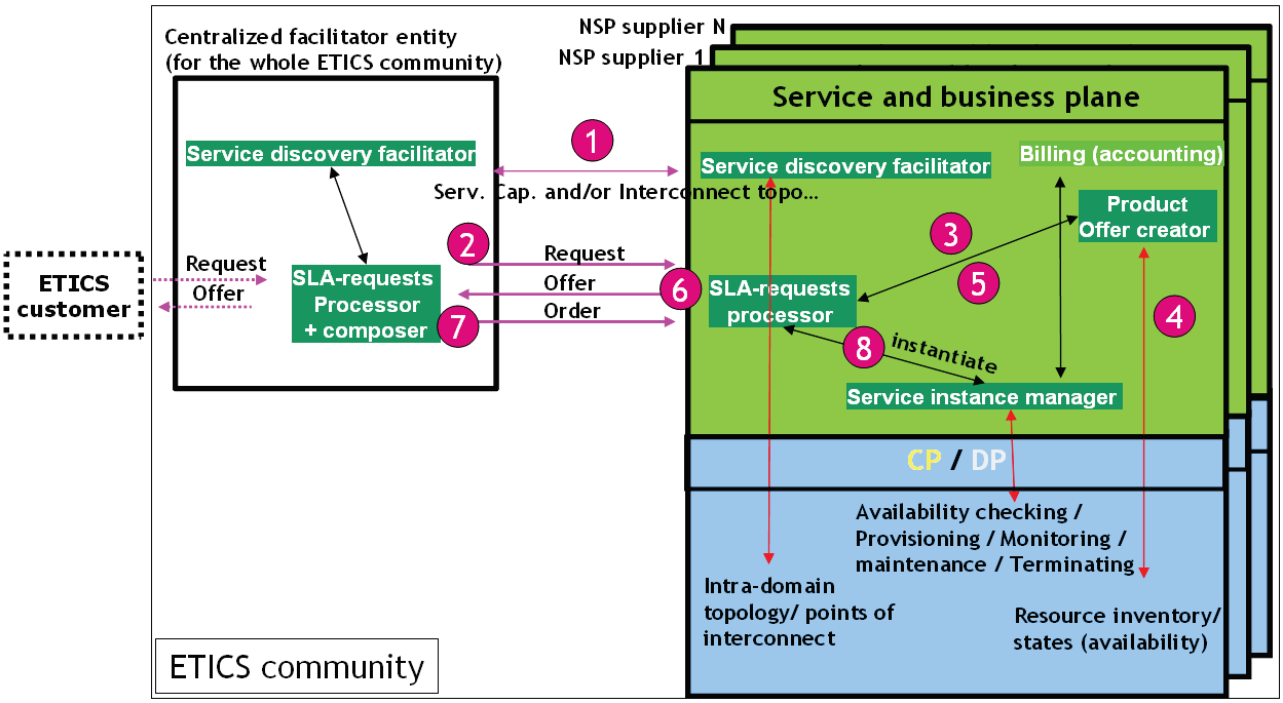


FIGURE 45: ON-DEMAND (PULL) SCENARIO WITH A CENTRALIZED ENTITY

The centralized entity has mainly two subentities. The first is the service discovery facilitator. Its role is to interact with the different service discovery facilitators at different NSPs and collect information that is

needed for NSP chain learning (called earlier Inter carrier topology or service capabilities). The second entity is the SLA requests processor and composer. This entity has two roles: the first is to process (receive/send) requests from external ETICS customers and to specific intra ETICS NSPs as customers. Its second role is the service composition. In a first step, it relies on the information gathered by the service discovery facilitator to potentially form one or more NSP chains that are probably able to provide a given inter-carrier product. In a second step, it sends specific requests to NSPs belonging to an NSP chain. This entity checks if the sum of the offers received by NSPs in an NSP chain meets the requirements of the inter-carrier product it wanted to form.

This scenario works as follows: In a first step, the centralized facilitator has interacted with all the service discovery facilitators in the ETICS community (1). This would have allowed the central entity to learn which NSP chain is the best suitable for a given request it receives. Upon the reception of an ETICS customer request, the central entity will send requests to all the NSPs in the community (2). An NSP receives this request thanks to the « SLA-request processor » entity which will trigger the offer creation (3). Once the offer is created (4&5), the SLA-requests processor proposes it to the centralized entity. If this offer is considered suitable, the central entity sends an order for it. Finally upon the reception of an order, the SLA-requests processor entity triggers the service instance manager to instantiate the offer. A first detailed version of this process is provided in the next section where we draw the sequence diagram involving the different actors of this feature.

#### 4.2.1.2. Inter-NSP Sequence diagram

FIGURE 46 shows the inter-NSP diagram relative to this first scenario. In a first step, NSPs independently send their network *capabilities* to the centralized facilitator. This step allows the centralized facilitator to compute a set of AS paths/NSP chains that correspond to an end-to-end connectivity request.

Upon the reception of a customer request, the centralized facilitator first computes one or more *NSP chains* (e.g. [NSP1…. NSPn]) utilizing the capabilities it had already received from the different NSPs.

In this scenario, we assume that the network capabilities are enough detailed to allow the centralized facilitator to split the QoS budget among the different NSPs in the chain. Once done, it then separately sends requests for specific QoS to each of the NSPs. Each NSP answers by either providing an offer that corresponds to the facilitator's request, or by signalling that it cannot satisfy the offer. If the centralized facilitator can combine the offers to obtain an end-to-end one that satisfies the customer request, it proposes it to the customer. If the customer is satisfied with the offer, it can order it from the facilitator. The facilitator will then send order requests to each of the NSPs.

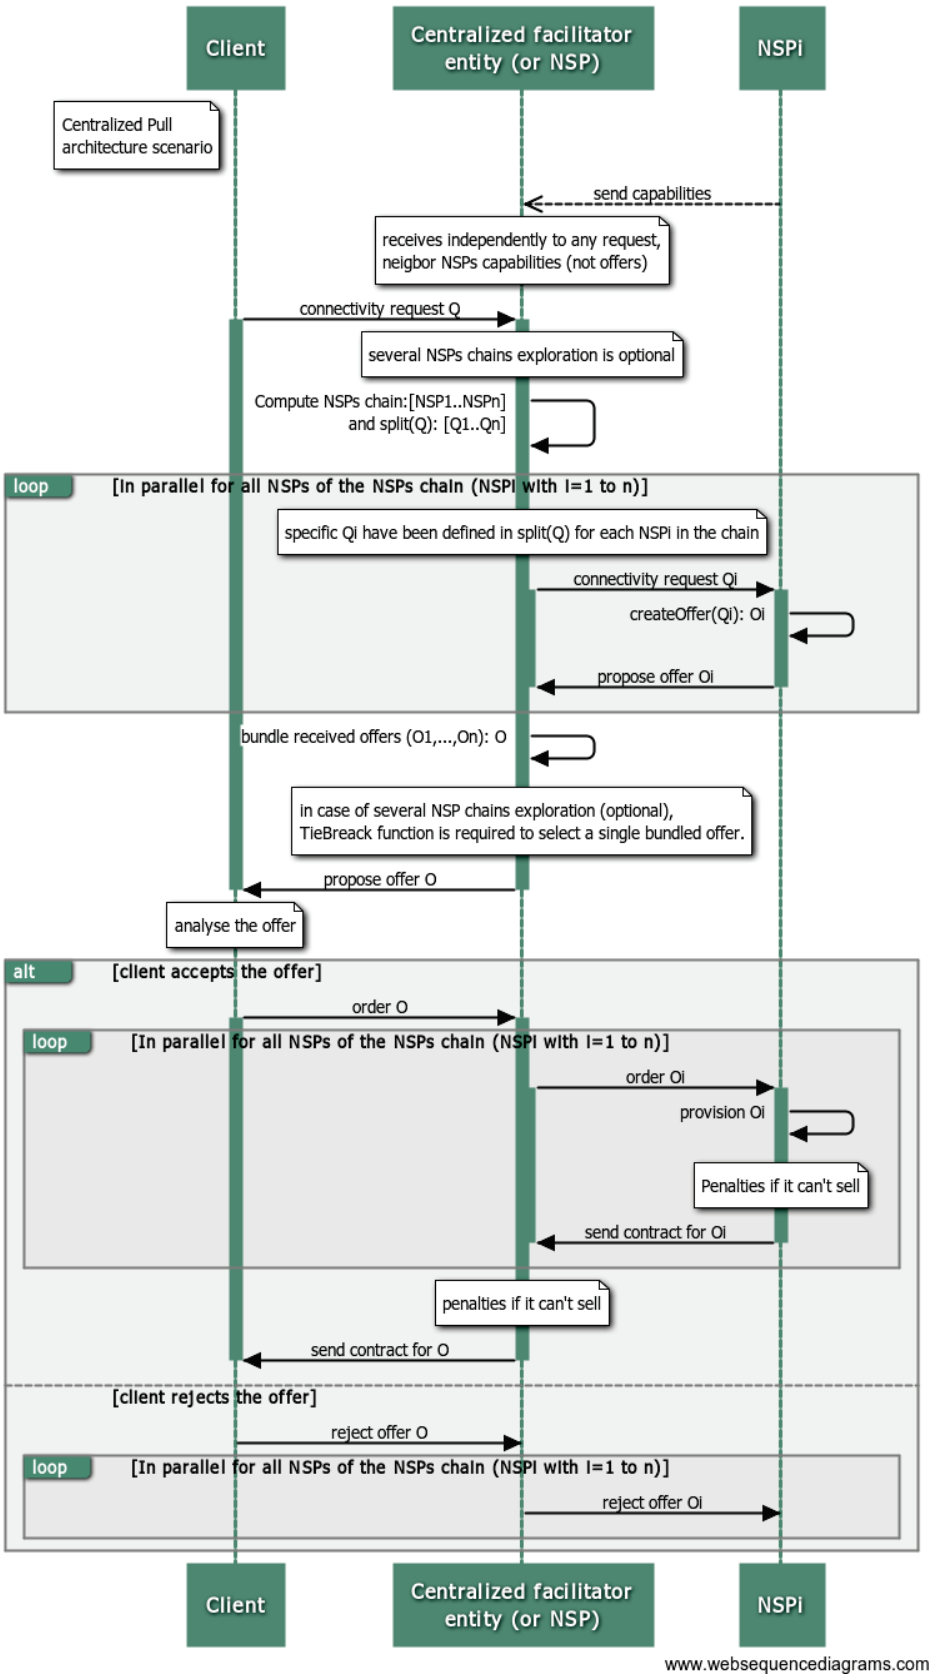


FIGURE 46: ON-DEMAND (PULL) SCENARIO WITH A CENTRALIZED ENTITY (INTER-NSP SEQUENCE DIAGRAM)

#### 4.2.1.3. Intra-NSP Sequence diagram (D4.2 terminology)

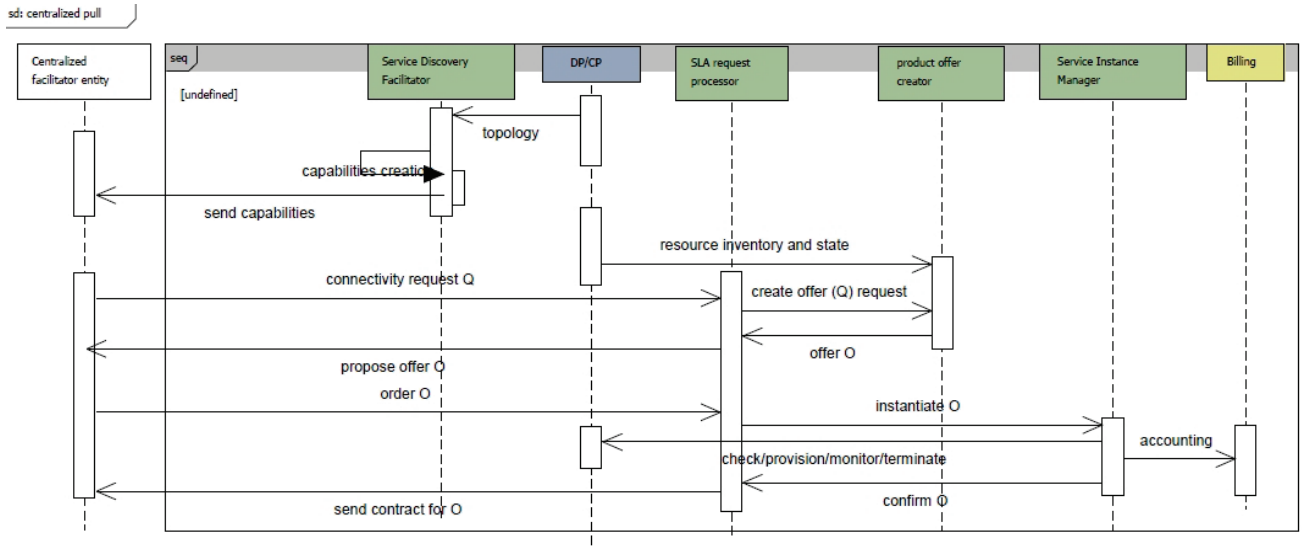


FIGURE 47: ON-DEMAND (PULL) SCENARIO WITH A CENTRALIZED ENTITY (INTRA-NSP SEQUENCE DIAGRAM)

To better illustrate the functioning of the on-demand centralized feature, we show in FIGURE 47 how the different intra-NSP building blocks interact with the centralized facilitator entity.

#### 4.2.1.4. Implications on the implementation

This section identifies the blocks that need to be specified and implemented for the centralized on-demand feature. First, the service discovery needs mainly two mechanisms:

- **Service capabilities exchange/IC routing protocol:** A protocol/mechanism to make the central facilitator entity aware of all service capabilities. This knowledge should be up to date (updated upon any change of service capabilities). The mechanism can be as simple as a repository managed by the central facilitator entity and updated by the different NSPs upon service capabilities changes. This protocol is called the IC routing protocol in the SLA class diagram of FIGURE 27.
- **NSP chain computation:** A system and/or an algorithm that takes as an input the information of the inter-carrier topology database. Upon the reception of an end-to-end ASQ path request, it computes the NSP chain(s) together with the corresponding capabilities that could satisfy the end-to-end request.

Second, the service composition step needs mainly:

- **Splitting the end-to-end QoS budget:** Assuming capabilities are enough detailed to allow for budget splitting between the different NSPs, the centralized facilitator is responsible for splitting the QoS budget between the different NSPs. A QoS splitting algorithm is needed to perform this task. The algorithm must be fair and different NSPs in the ETICS community must agree on it.
- **SLA requests exchange/IC signalling protocol:** A negotiation protocol to exchange SLA requests and offers between the centralized facilitator and the different NSPs in the ETICS community.
- **Associate capabilities and offers with the "network layer":** If the offers are not enough detailed (lack of the exact router/interface that will be used for traffic delivery), then we need a mechanism

to ensure that the provisioned path is the same as the offered one. This is currently a work in progress.

- **Provisioning an end-to-end path:** This task consists in provisioning the end-to-end path. The way the path is provisioned depends on the different technologies that each NSP supports.

### 4.2.2. ON-DEMAND (PULL) OFFER WITH PER-NSP CENTRALIZED COMPOSITION SCENARIO

#### 4.2.2.1. Architecture- High Level View

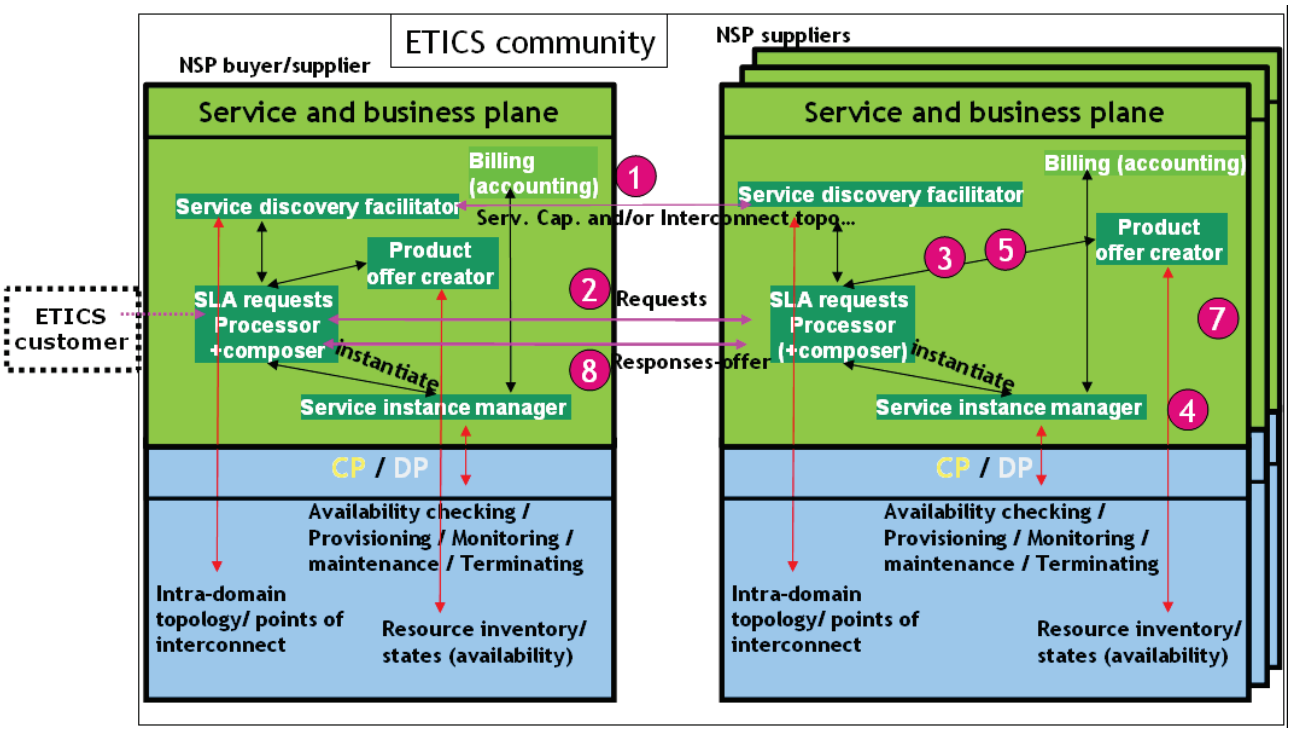


FIGURE 48: ON-DEMAND (PULL) SCENARIO WITH PER-NSP CENTRALIZED COMPOSITION

FIGURE 48 picks up the ETICS service and business plane architecture in the case of on-demand scenario from [ETICS-D4.2], with distributed NSP chain learning and a centralized service composition ***per NSP***.

We remind that with respect to the previous case of the on-demand variant with centralised composition entity (cf. Section 4.2.1), it is as if the centralized entity's functions were implemented in a distributed way at the level of each NSP. Each NSP can therefore play both roles. The first is the role of a central entity that will make service composition, and secondly, the role of a supplier that will answer upon the reception of a request from an NSP playing a central composition role. It is important to note that in this feature, the NSP that plays the central role **is in charge of splitting the end-to-end QoS budget** among the different NSPs in an NSP chain.

#### 4.2.2.2. Inter-NSP sequence diagram

The inter-NSP sequence diagram is the same as that of Figure 46 except that the centralized facilitator role can be played by any NSP. NSPs independently exchange network capabilities between each other. A

customer can request an end-to-end ASQ path from any NSP. The NSP receiving the service request plays the central facilitator role to compose the service and order the offers.

**We also assume for this scenario that the service capabilities are sufficiently detailed to allow each NSP to split the end-to-end QoS budget between the different NSPs in an NSP chain**.

#### 4.2.2.3. Implications on the implementation

The per-NSP centralized scenario is very similar to the one with a central facilitator entity. The only difference is that the central facilitator entity functions need to be implemented in the service and business plane of each NSP. The implications on the implementation described in Section 4.2.1.4 are the same as for this scenario:

- **Service capabilities exchange/IC routing protocol:** A protocol/mechanism to flood service capabilities and to update them upon a change. This could be done utilizing an overlay of per NSP (or Autonomous System) OSPF routers. This solution, which is currently being specified within ETICS, is called H-TE (Hierarchical Traffic Engineering). In this solution, the reliable flooding mechanism of OSPF is responsible for the flooding of capability information between the different routers. Other possibilities include the use of web services.
- **NSP chain computation:** An algorithm that takes as an input the information of the inter-carrier topology database. Upon the reception of an end-to-end ASQ path request, it computes the NSP chain(s) together with the corresponding capabilities that could satisfy the end-to-end request.
- **Splitting the end-to-end QoS budget:** This can be possibly done in different ways depending on the level of detail concerning the service capabilities. If the capabilities are enough detailed to allow for budget splitting between the different NSPs, the centralized facilitator is responsible for splitting the QoS budget between the different NSPs. A QoS splitting algorithm is needed to perform this task.
- **SLA requests exchange:** A negotiation protocol to exchange SLA requests and offers between the centralized facilitator and the different NSPs in the ETICS community.
- **Associate capabilities and offers with the "network layer":** If the offers are not detailed enough (lack of the exact router/interface that will be used for traffic delivery), we need a mechanism to ensure that the provisioned path is compliant with the offered one (e.g. it goes through the same points of interconnect).
- **Provisioning an end-to-end path:** This task consists in provisioning the end-to-end path. The way the path is provisioned depends on the different technologies that each NSP supports.

### 4.2.3. On-demand (pull) Offer with Distributed Composition Scenario

#### 4.2.3.1. Architecture- High Level View

FIGURE 49 shows the ETICS service and business plane architecture in the case of on-demand scenario, with distributed NSP chain learning and distributed service composition, as defined in [ETICS-D4.2].

In this scenario, **as a first step, different NSPs exchange more or less detailed intra-NSP information called service capabilities** (steps 1 and 1'). It is important to notice that, unlike the first two features where the

service capabilities need to be detailed enough to allow for QoS budget splitting, in this scenario their role is only to help finding the best NSP chain(s). As such, at the end of his step, each NSP is able to know which NSP chains are more likely to provide the service given a particular ETICS end request.

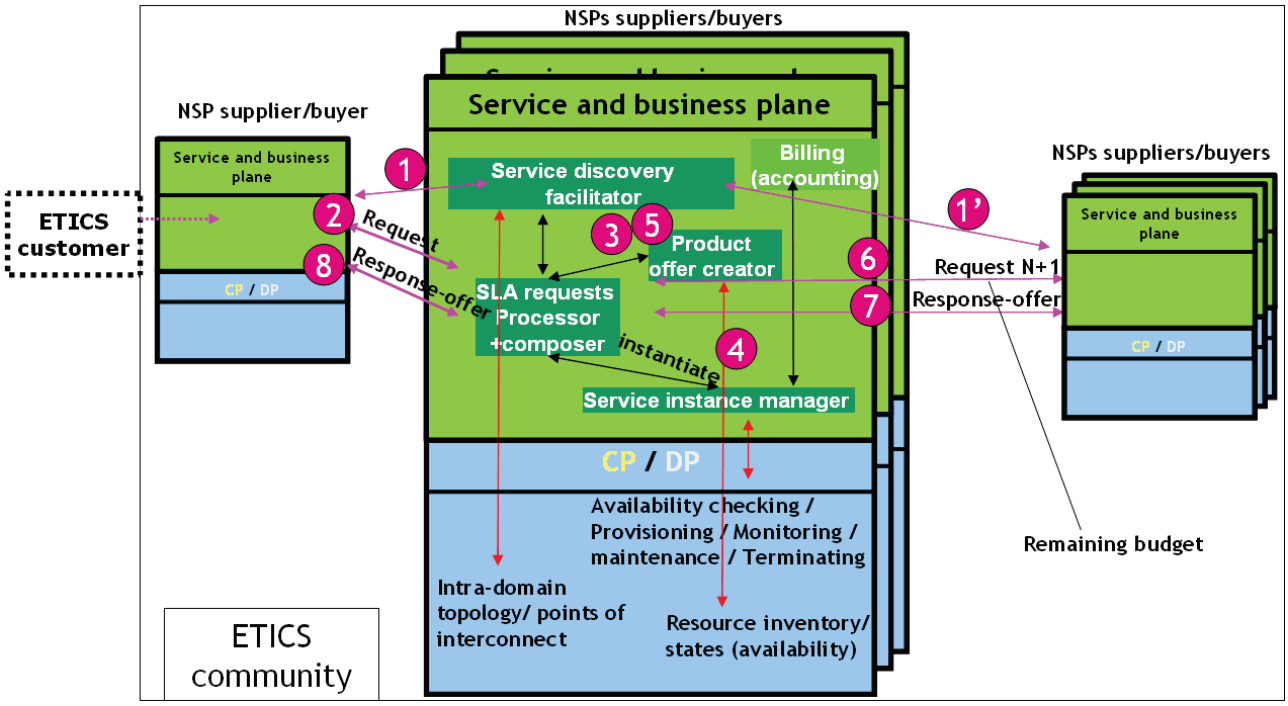


FIGURE 49: ON-DEMAND (PULL) SCENARIO WITH distributed composition

In a second step, the service composition will be done in a decentralized fashion, following one of the cascading models described in prior deliverables. **The splitting of the QoS budget will be done in a distributed way**. Each NSP decides its own "contribution" to the overall end-to-end QoS budget.

As mentioned in D4.2 [ETICS-D4.2], the «SLA requests processor» entities of NSPs receive SLA requests from neighbouring NSPs (2). These requests will trigger offer creation (4, 5). A new request, with a remaining budget is then sent to the next NSP(s) in the chain(s) for a forward recursive service computation strategy. For instance, if the received request contains 100 ms as the delay value for reaching a destination and if the receiving NSP will consume 20 ms in order to reach the next NSP in the NSP chain, then the next forwarded request will be 80 ms to reach the destination. If the offer composition succeeds and arrives to the destination, a set of cascading responses is then sent backward to the source entity which initiated the offer.

A similar backward recursive process can also be used (not presented in the figure). Such a scheme is very similar to the backward recursive path computation (BRPC) process [RFC5441]. In this case, the request is transmitted with no budget modification along a NSPs chain to the destination NSP. **The budget is then consumed when offers are relayed in a backward way**. No matter what composition strategy is used, forward or backward, the overall QoS budget relative to an ETICS customer request is split in a distributed way. Each NSP decides how much it would consume from the overall budget. This schema will evidently give more freedom (and possibly an advantage) to the first NSPs in the chain.

In the former forward strategy (described by FIGURE 49), once a final offer is chosen by the final ETICS customer, a second chain of order requests is then sent to instantiate the offers. This chain is not represented in the figure. Ordering an offer can be done, as the composition, in hop-by-hop cascading manner. It can also be done separately to each NSP in the chain. This depends more on the business agreements between the different NSPs in the ETICS community.

#### 4.2.3.2. Inter-NSP sequence diagram

FIGURE 50 shows the sequence diagram of this scenario. The main difference to the previous scenarios is that the service composition and specially the end-to-end QoS budget splitting are done in a distributed way. Each NSP that receives a request with a given QoS budget computes its own budget and forwards a new request to the next NSP with the remaining budget.

In the first two scenarios, service capabilities need to be more or less detailed to allow the central entity (facilitator or an NSP) to split the end-to-end QoS budget among the different NSPs in the chain. This scenario is slightly different in the way that service capabilities can be less precise. Their unique role is to help NSPs compute the chain of NSPs to request in order to satisfy a given inter-carrier request.

There are possibly other variations of this scenario, which could result from the use of already existing standards like the Path Computation Element (PCE) architecture. In a first step, PCE can be used for the exchange of service capabilities needed in order to compute the NSP chain. A standard BRPC procedure would then be triggered in a second step in order to compute an inter-NSP path.

#### 4.2.3.3. Implications on the implementation

In the distributed on-demand scenario, the service composition and especially the end-to-end QoS budget splitting is done in a distributed way. The implications relative to this scenario are described as follows (some of them were already described in Section 4.2.2.3):

- **Service capabilities exchange/IC routing protocol:** A protocol/mechanism to flood service capabilities and to update them upon a change. This mechanism is the same as the one described in Section 4.2.2.3.
- **NSP chain computation:** An algorithm that takes as an input the information of the inter-carrier topology database. Upon the reception of an end-to-end ASQ path request, it computes the NSP chain(s) together with the corresponding capabilities that could satisfy the end-to-end request.
- **SLA requests exchange:** A negotiation protocol to exchange SLA requests and offers between the different NSPs in the ETICS community.
- **Ensure that capabilities/offers match with the "network layer":** If the offers are not enough detailed (e.g. lack of the exact router/interface that will be used for traffic delivery), we need a mechanism to ensure that the provisioned path is compliant with the offered one). This is ongoing work. Some first solutions have been designed within the consortium to tackle this problem.
- **A mechanism to concatenate the different offers together:** At provisioning time, a mechanism to stitch the different offers together is needed especially when no end-to-end control plane signalling (e.g. RSVP-TE) is possible, or when the path crosses heterogeneous domains.

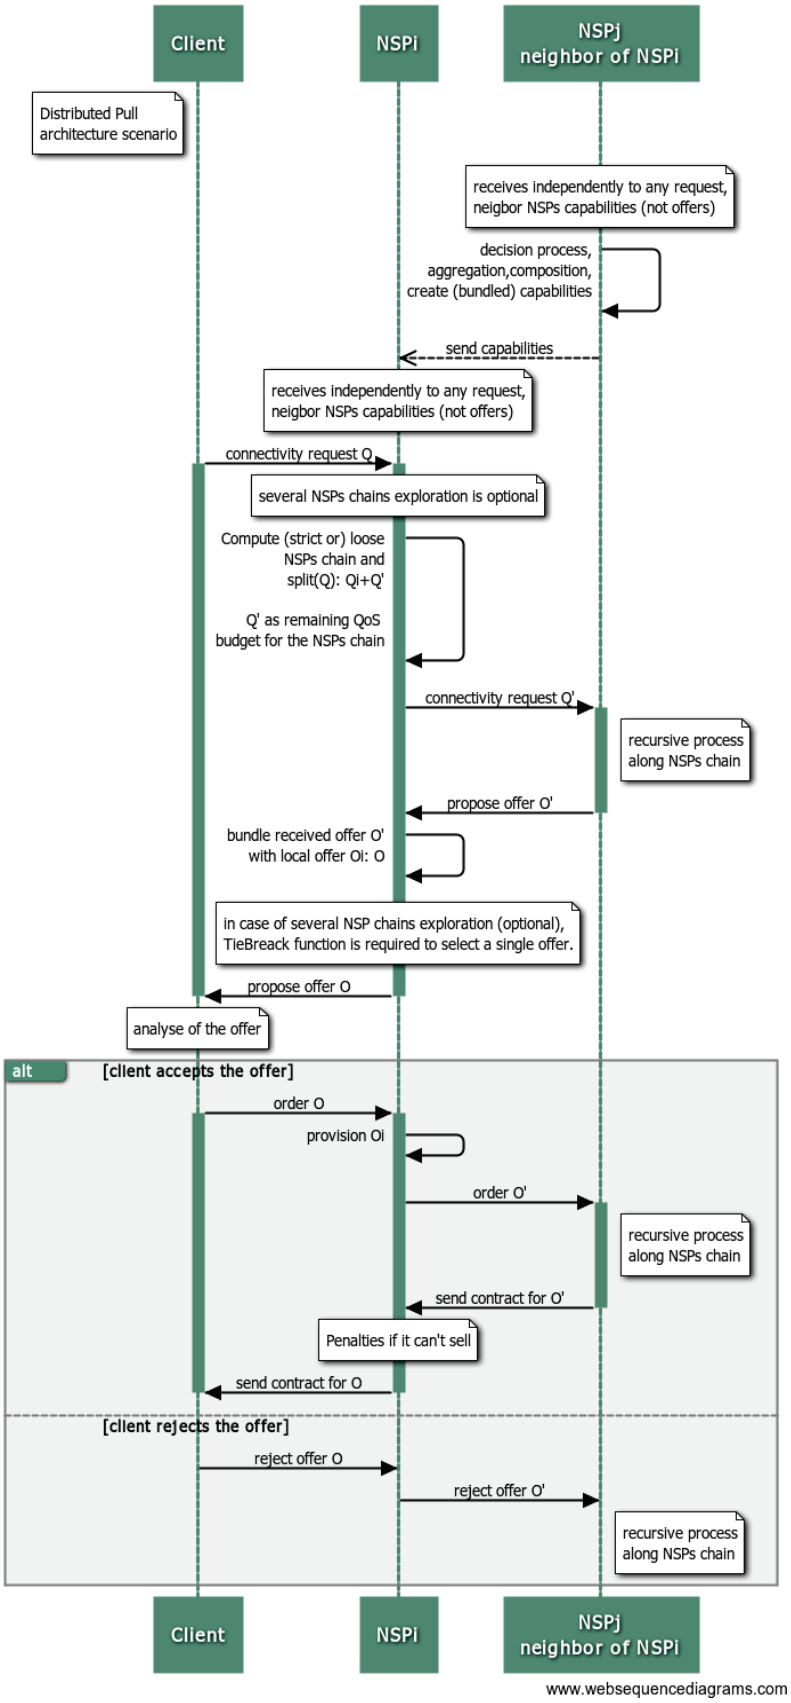


FIGURE 50: DISTRIBUTED ON-DEMAND (PULL) SCENARIO (INTER-NSP SEQUENCE DIAGRAM)

### 4.2.4. PRE-COMPUTED (PUSH) OFFER WITH UNIQUE CENTRALIZED COMPOSITION ENTITY SCENARIO

#### 4.2.4.1. Architecture High-level View

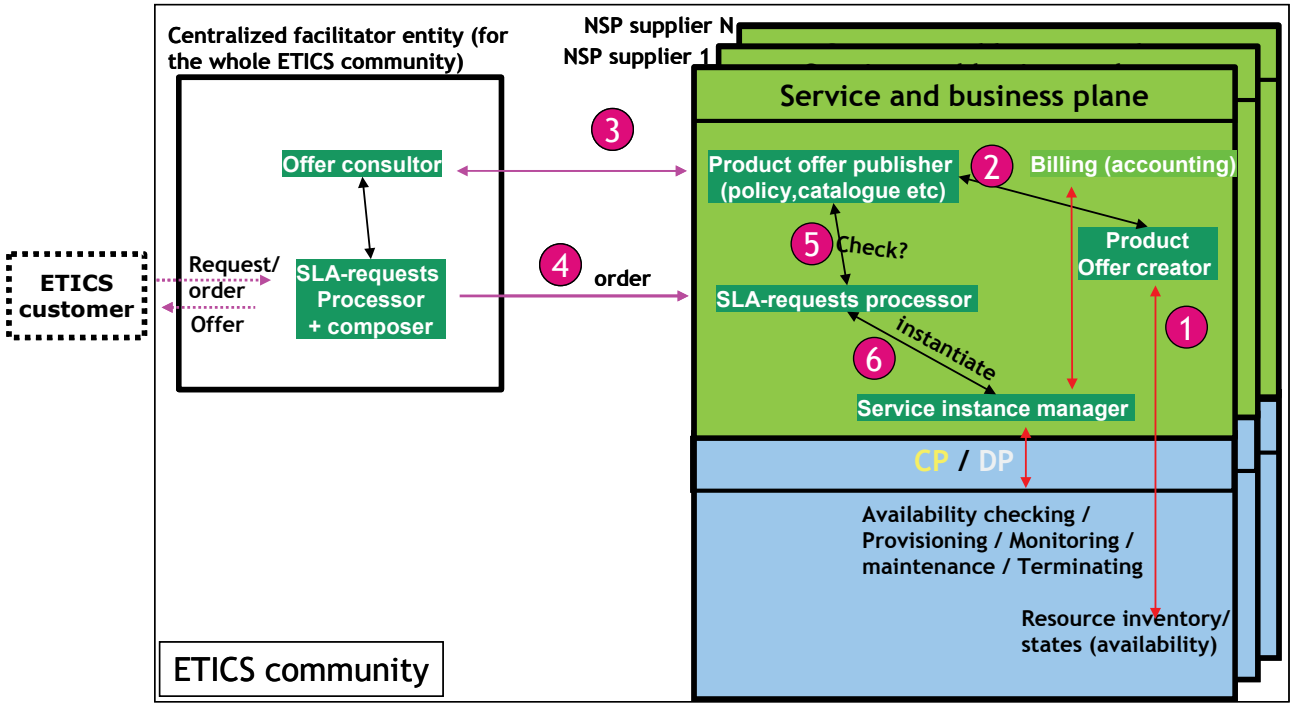


FIGURE 51: PRE-COMPUTED (PUSH) SCENARIO WITH A CENTRALIZED ENTITY

FIGURE 51 shows the ETICS service and business plane architecture in the case of the "push" scenario with a centralized third party entity that handles the service composition process, as presented in deliverable D4.2. In this scenario, offers are first pre-computed by NSPs, and are then made accessible to the centralized facilitator entity, which uses these offers to compute a global inter-carrier offer. In this scenario, the service composition is done by the central entity upon the reception of an ETICS customer request for an end-to-end inter-carrier ASQ path. The central entity, which has access to all the individual per-NSP offers, needs "only" to compute the best combination that can fit the customer request.

For this feature (as described in D4.2), offers are first (1) pre-computed by the product offer creator entity (called SLA offers builder in the recent UML class diagram), and are (2) handled to a second entity, the product offer publisher (renamed SLA offers protocol in the UML class diagram).

The centralized entity has mainly two sub-entities. The first one is the offer consulter (updated in the UML class diagram to the SLA offers protocol). The central entity uses this interface or protocol to access all the offers of the ETICS community. The second entity is the SLA requests processor and composer. This entity has two roles: It handles the requests for inter carrier products (either from intra ETICS NSP customers or external ETICS customers), while its second role is to compose offers. Upon the reception of an inter-carrier product request for instance, this entity computes which intra-ETICS offers can be concatenated together to meet the inter-carrier ASQ path request. It can also order an offer (4) which would trigger offer instantiation (5).

Since the SLA request processor and composer have two roles, they have been revised in the newer version of the architecture (Sections 4.1.2 and 4.1.3.3) by splitting it into two functional blocks: the IC signalling protocol for SLA request processing and exchange, and the SLA controller for service composition.

#### 4.2.4.2. Inter-NSP sequence diagram

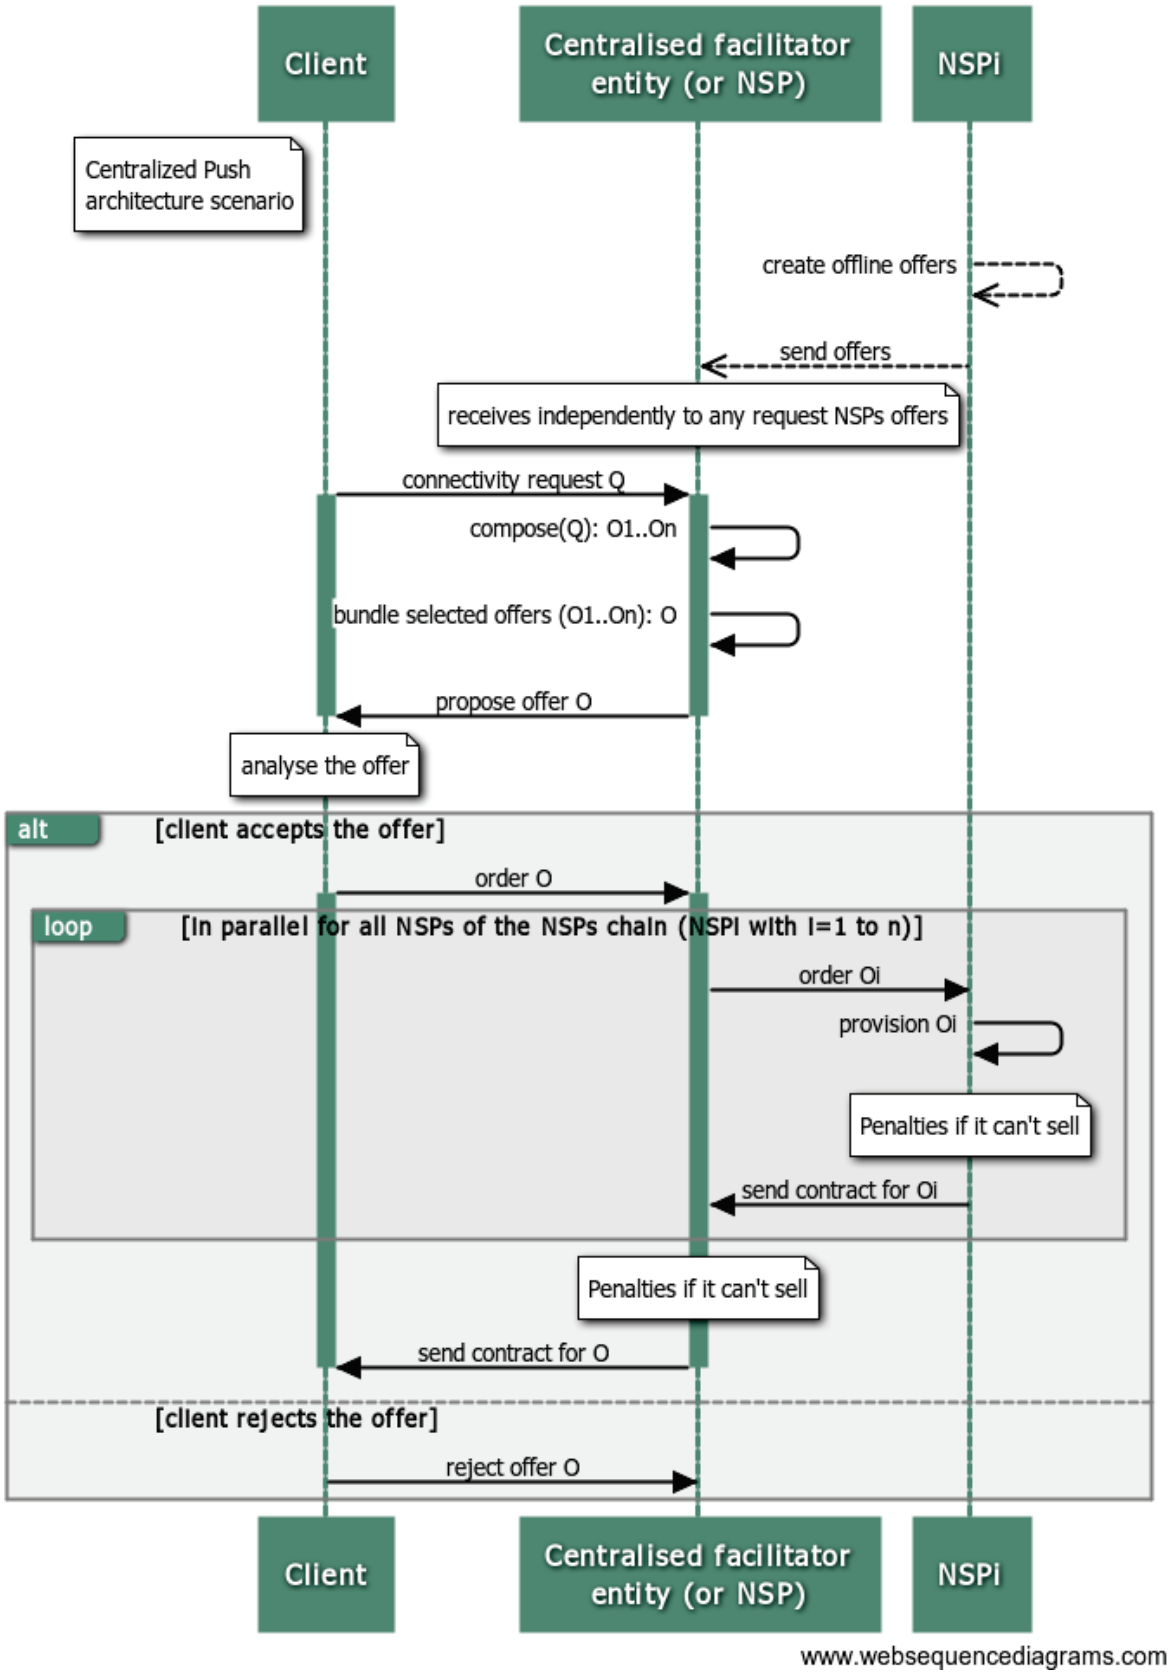


FIGURE 52: PRE-COMPUTED PUSH SCENARIO WITH A CENTRALIZED ENTITY (INTER-NSP SEQUENCE DIAGRAM)

FIGURE 52 shows the inter-NSP (NSP granularity) sequence diagram of the centralized push scenario. The sequence diagram is valid for both, the centralized facilitator entity, as well as the per-NSP centralized case. The figure shows three actors: the ETICS client, the central entity (NSP or facilitator), and neighbouring NSPs.

In a first step, NSPs create their offers. Then, they either send them to the facilitator or they exchange them between each other. As such, each central entity has the offers of all the NSPs in the community. Upon the reception of an ETICS client request for offer, the central entity computes the best set of offers that can satisfy the customer request. The central entity sends the best end-to-end offers to the ETICS client. The client can either pick one or refuse all the offers.

Depending on whether the offers are enough detailed, i.e. can be provisioned directly because they contain network information (router, interface etc.) or not, an extra step might be needed to compute the exact "deployable" path that should be implemented in the network. If the networks are heterogeneous, a further step might be needed to allow for "stitching" the different per-NSP offers together.

#### 4.2.4.3. Implications on the implementation

In the centralized push scenario, the service composition is done by computing the chain of offers that satisfy a given ASQ path request from a customer. The implications on the implementation relative to this scenario are the following:

- **Offers exchange/SLA offers protocol:** A protocol/mechanism that allows NSPs to send their offers to the centralized facilitator and to update them upon a change. The centralized facilitator keeps a repository that can be accessed by the NSPs to add, update or remove their own offers. This protocol is similar to the IC routing protocol.
- **End-to-end offer computation:** An algorithm that takes as an input the offers catalogue. Upon the reception of an end-to-end ASQ path request, it computes the set of offers that could satisfy the end-to-end request. This algorithm is very similar to the NSP chain computation algorithm since offers and service capabilities have similar structure (ingress point/point of interconnect, egress point, a QoS vector, etc.)
- **SLA requests exchange/IC signalling protocol:** A negotiation protocol to exchange SLA requests (e.g. order an offer) between the centralized facilitator and the different NSPs in the ETICS community. This protocol is very similar to the one that is used in the on-demand scenarios.

If the offers are not enough detailed to allow for the ordering and the stitching of offers at the network (data plane) layer, extra mechanisms might be needed:

- **Ensure that offers match with the "network layer":** If the offers are not enough detailed (lack of the exact router/interface that will be used for traffic delivery), then a mechanism to ensure that the provisioned path is the same, as the offered one is needed. Note that this is work in progress.
- **A mechanism to stitch the different offers together:** At provisioning time, a mechanism to stitch the different offers together is needed especially when no end-to-end control plane signalling (e.g. RSVP-TE) is possible (e.g. when the path crosses heterogeneous domains).

### 4.2.5. PRE-COMPUTED (PUSH) OFFER WITH PER-NSP CENTRALIZED COMPOSITION SCENARIO

Similarly to case "ON-DEMAND (PULL) OFFER WITH PER-NSP CENTRALIZED COMPOSITION SCENARIO", depicted in section 4.2.2, the centralized entity role can be implemented at the level of each NSP. In this case, each NSP can play the role of the central entity. We do not detail this case as it is similar to the central entity case of Section 4.2.4.

### 4.2.6. PRE-COMPUTED (PUSH) OFFER WITH DISTRIBUTED COMPOSITION SCENARIO

#### 4.2.6.1. Architecture High-level View

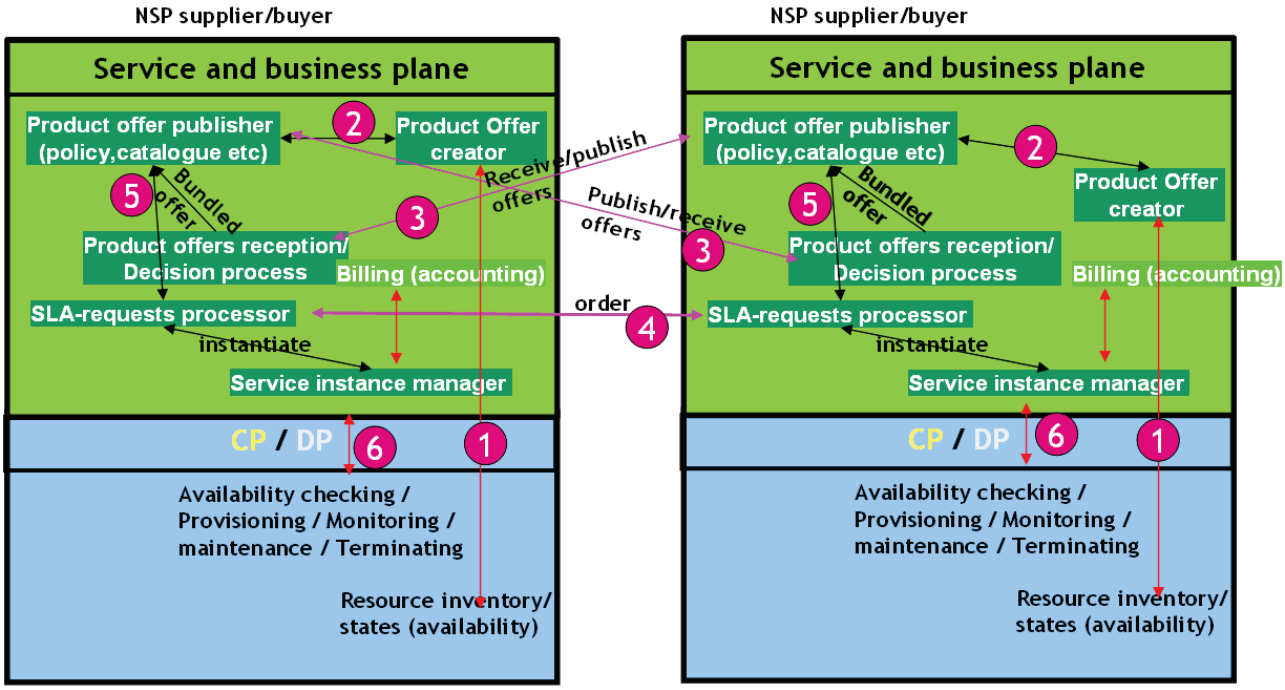


FIGURE 53: PRE-COMPUTED (PUSH) SCENARIO WITH DISTRIBUTED COMPOSITION

FIGURE 53 presents the ETICS service and business plane architecture in the case of a push scenario and a distributed service composition process, as presented in D4.2. In this scenario, offers are first created and then published to NSPs in the community. An NSP that gets aware of an offer can either order it for its own use, or aggregate it with its own offers and republish the aggregated offer. In this case, the service composition is done utilizing the propagation of aggregated offers.

This feature is very similar to the way the Border Gateway Protocol (BGP) operates: Upon the reception by a router of a route to a destination, a new route to the destination is formed and is then propagated. Similarly to BGP's decision process, an offers decision process is needed in order to choose between different offers concerning the same destination.

Compared to previous scenarios, in the D4.2 description a new entity was introduced, the «Product offers reception/decision process». Its role, as defined in D4.2, is:

- «Product offers reception/decision process»: This entity has two roles. The first is to compare different received offers. In fact, an NSP can receive different offers to reach the same destination

but with different SLSs. The second role is to bundle received offers from other NSPs with intra NSP offers to create new offers that will be republished again thanks to the product offer publisher.

In this deliverable, the scenarios defined in D4.2 are currently seen as features that work towards the same goal: providing inter-carrier ASQ paths. They can be as such integrated within the same overall architecture. With respect to the SLA class diagram presented in Section 4.1.3.3, the "product offers reception/decision process" is split into two blocks. The first is the *SLA offers protocol*, which is the interface of the NSPs to exchange offers between each other. The second is the *SLA offers controller* which has methods to compare and bundle offers etc.

#### 4.2.6.2. Inter-NSP sequence diagram

FIGURE 54 shows the inter-NSP sequence diagram relative to the distributed pre-computed scenario. In this scenario, each NSP creates its own offers. Depending on its policies, the NSP propagates its offers in the form of a vector to (some of) his neighbours. An NSP that receives an offer applies its import policies to see if it imports it (interested in it). Then, the NSP has the choice between keeping the offer for itself or aggregating it with its own ones and sending it to its neighbours. Proceeding this way, each NSP will have a set of end-to-end offers involving all the NSPs in the chain.

#### 4.2.6.3. Implications on the implementation

The blocks that need to be specified and implemented for this feature are:

- **A protocol to exchange offers:** Similarly to BGP, the Border Gateway Protocol, this protocol builds a signalling graph between the different NSPs. This graph will be used to propagate the NSPs offers.
- **Offers format definition:** Offers might need to carry attributes that help avoiding loops or that specify policies. These attributes can be intended to influence the decision of other NSPs.
- **A decision process:** This process will influence the way offers are propagated inside the network. This process should avoid creating loops. It can also be designed to allow for the offer computation step to converge. Note that similar problems were detected with BGP when the routes propagation graph did not have a given pattern [GaRe01, GrSh02].

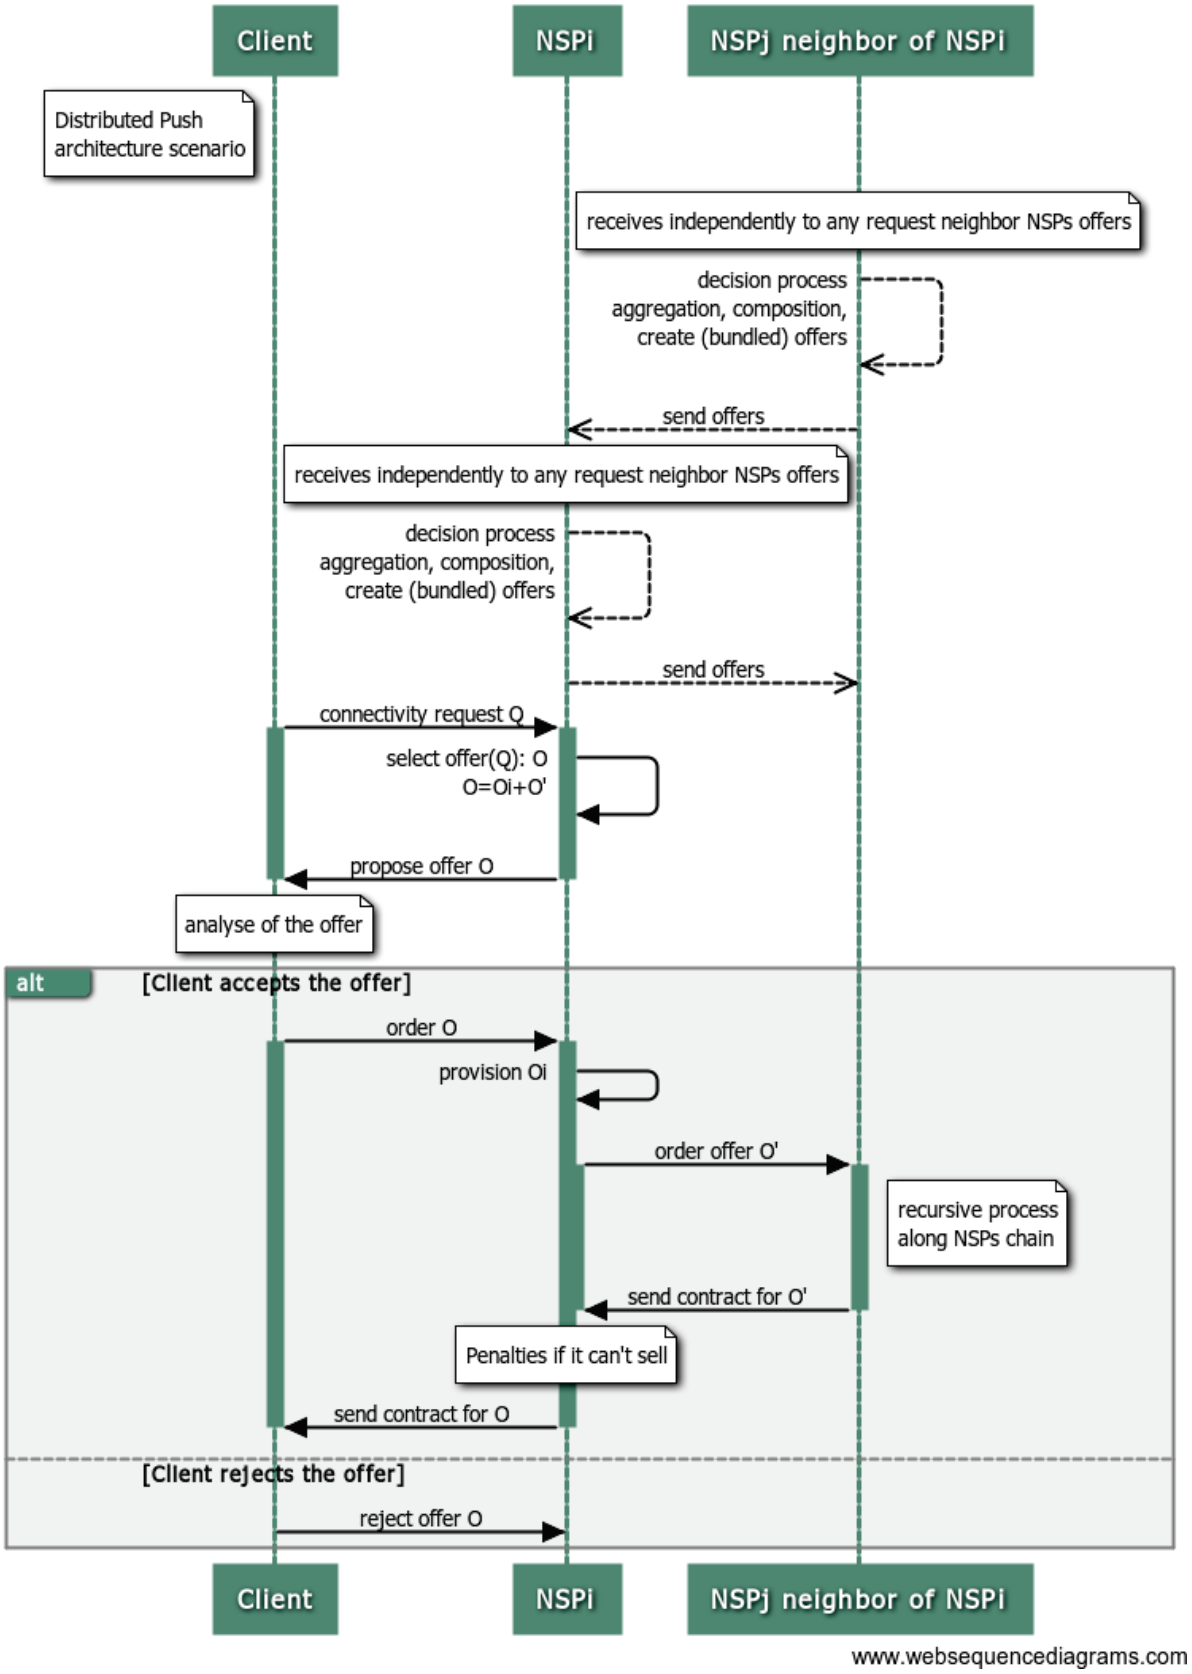


FIGURE 54: PRE-COMPUTED (PUSH) SCENARIO WITH DISTRIBUTED COMPOSITION (INTER-NSP SEQUENCE DIAGRAM)

# 5. PRELIMINARY PERFORMANCE AND SCALABILITY ANALYSIS

## 5.1. INTRODUCTION

This section collects information from WP3 and WP5 in order to provide a feedback on the evolution of the ETICS architecture from D4.2.

Firstly, a study of scalability has been performed, in which different aspects of the architecture that could be affected by scalability issues have been analysed. The target has been to review the *push* and *pull* models, most specifically related with the *Service and Business Plane*, and to underline which solutions are most scalable.

Economic concepts have been employed by WP3 in order to provide feedback on the architecture from the business and economic point of view. Accordingly, an assessment of different architectural options has been made through a set of economic criteria.

Finally, a performance analysis with partial simulation results has been executed by WP5 in order to assess different aspects of the ETICS Architecture. The performance analysis by WP5 has experienced some internal delays in recent months, such that some simulations in the present document unfortunately do not provide a too detailed assessment. Therefore, this work will be complemented by additional simulations in 2012, and the corresponding results will be integrated into D4.4 as the next iteration in the series of architecture documents.

## 5.2. PRELIMINARY ASSESSMENT OF THE SCALABILITY OF THE SCENARIOS

In this subsection, push and pull architectural scenarios (SB plane) are elaborated together with possible scalable solutions. All the functionalities and elements of the architecture have been confirmed and specific solutions have been elaborated in order to allow the system to scale, mainly concerning CPU, storage, queuing, proxies, etc.

### 5.2.1. SCALABILITY OF THE CONTROL PLANE: A SHARED CP CAN BE A BOTTLENECK

The scale of the control plane ultimately determines the ability to scale and operate independent services. Providers need the highest scaling parameters for each individual service, with minimal scaling collisions and secure isolation (error containment) between the services.

Even with the separation of the control and forwarding planes in individual network elements, the addition of new services to a single control plane can cause issues with stability, scale, and the processing requirements (see FIGURE 55).

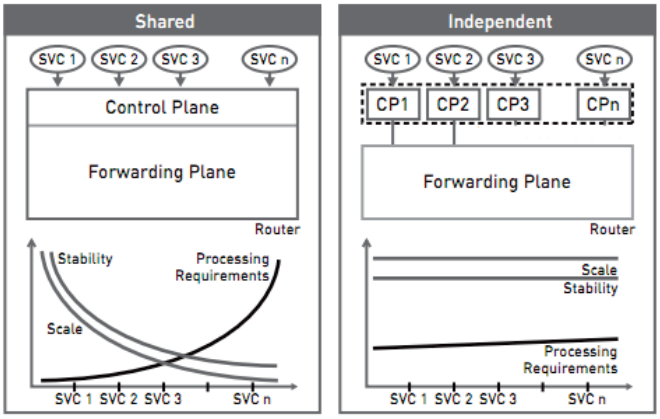


FIGURE 55: SHARED VERSUS INDEPENDENT CONTROL PLANE

For example, if a network element supports two or three business services, adding a consumer service means that an entirely new set of compound scalability tests must be run on the network element.

However, if the control plane can be partitioned into individual planes supporting separate services, then the scale, stability, and processing requirements become much more predictable.

This is an environment in which new services can be rapidly prototyped and much more smoothly rolled out. It facilitates flexible, reduced-risk service enablement that translates into rapid introduction of new services while reducing complex testing and capacity planning.

It also results in greater administrative separation, allowing different business groups within a service provider to administer different service instances.

Beside that, the scalability of the control plane has to satisfy the requirement TR-NET-PC-06 as recalled below. For this purpose, subsequent paragraphs will be dedicated to this analysis.

**TR-NET-PC-06:** "The path computation method must be able to manage an increase of the following contextual parameters: number of carriers, number of path computation requests."

Intra carrier

The scalability should take into account the load balancing of PCE based on internal domains of the carrier. It means that the dimension of the carrier network must be taken into account when scaling up the PCE / control plane. The TED (Traffic Engineering Database) is also impacted and needs to be scaled up.

The computational part (CPU) has to scale with the dynamic evolution of the network topology and PCE REQ / RESP (frequency).

Inter Carrier

The scalability process is comparable to the intra domain case that has to scale with the network topology inter carrier / meshed and its related TED.

The computational part (CPU) has to scale with the dynamic evolution of the network topology inter-carrier and PCE REQ / RESP (frequency).

### 5.2.2. PUBLISH SCENARIO (PUSH MODEL)

Such an approach may be exhaustive if all possibilities are announced, which certainly bears a scalability problem (e.g. if the validity time of the offers is too short or if the frequency of service composition is too high) when the number of NSPs becomes too large.

For example, we can make the following assumptions in a first assessment of this issue:

- An NSP AS (in particular transit domain) uses around 30 to 60 ASBRs,
- 4 to 8 offers per ASBR couples (or POI) are assumed (e.g. one per CoS).

This results in 3 000 to 60 000 offers which contain network prefixes, and for 60 ASes in the ETICS consortium this corresponds to 180 000 to 3 600 000 SLA offers.

Therefore, this scenario must be able to handle several millions of offers. This figure can be considered as huge when used in a given critical context (e.g. the Forwarding Information Base (FIB) routing table in routers), but it does not necessarily represent a big issue in other contexts, such as setting up IC ASQ paths of large capacity (N-SLA) between NSPs, where the information will be consulted a few times a day (e.g. 10000 per NSP), and where reasonable waiting times are allowed (e.g. 10 seconds[23]) in order to obtain the result of service composition.

In addition, if we just look at the originating prefix network announced by an NSP (around to 1000 to 10000 for a transit AS), the number of offers could strongly be reduced. Indeed, only offers for a given destination are of importance. Therefore, as a first approximation, we may just publish the different offers (e.g. one per Class of Service – CoS) for each destination network. This would allow for filtering among the huge number of offers to select which ASs could be selected in order to build a complete SLA end-to-end offer. Once the AS chain is selected, a second step in the process will refine the offer by selecting the TDP and thus the ASBR interfaces couple.

By introducing more precision to our initial assumptions, we obtain:

- An NSP transit announces from 1000 to 5000 originating network prefixes
- An Edge NSP announces from 100 to 1000 network prefixes
- Assuming 10 Transit and 50 Edge NSP, we get between 15 000 and 100 000 offers, and if combined with several Class of Service options (Q) per offer we could reach 60 000 to 400 000 SLA offers for 4 CoS.

This time, the number of SLA offers is in the same order of magnitude as the number BGP of routes. Simulations will be performed in WP5 which shall confirm and assess these figures more precisely in order to reaffirm the scalability of the approach.

In FIGURE 56 below, we show an example of how an NSP could compute its offer. This example relies on the usage of the control plane information to compute offers in a fully automated process: In a first step, the service plane collects from the control plane all couples {Q – quality, C – cost} for each couple of ASBRs

[23] These numbers are just used as indications but they will have to be more precisely defined in the future. Studies planned in the WP5 will provide for a clearer view on the different limits.

associated to a given destination prefix. Then, it determines the sets of time parameters T (time) from its Policy rules database.

The Business plane in step three parses all C (cost) parameters in order to compute the price P again, using the Policy rules database.

Finally, the ETICS system publishes a set of {N – network prefix, Q – quality, T – time, P – price} through the *E1*, *E2*, *E3* interfaces where N design the network prefix that could be reach with the QoS Q, for the time T and at price P.

In this step of the process, ASBR and TDP need not be known. Indeed, service composition has sufficient information to determine the AS chain. Once the AS chain is known, the validation step in the SLA life cycle will allow fixing the TDP, and thus, the ASBR couple. Offers could be announced to all neighbours or to a particular neighbouring AS. The policy rules database helps the business plane to adjust the offers regarding the peer ASes, and specialises the offers.

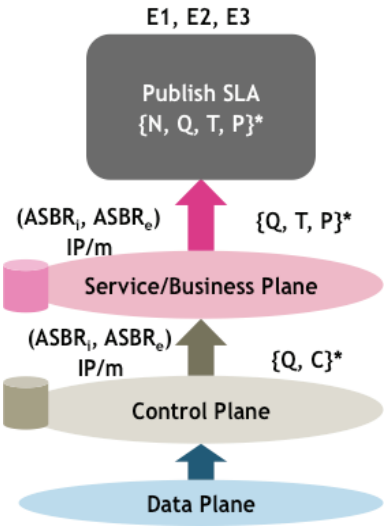


FIGURE 56: PUBLISH SCENARIO FOR SLA COMPOSITION

Even if the number of offers is large, the most crucial part is the synchronisation of offers between the different NSPs. Indeed, if SLAs are indefinitely available for the service composition process, they could be out of date. Time parameters are defined for that purpose, in order to let service composition discard all offers which have reached the defined timeout. However, an NSP could choose to update an offer when its timeout is reached in order to keep the offer valid. Again, this greatly depends of the frequency of the SLA updates done by the NSP.

Scalability of the scenario depends on the frequency of these SLA updates with implications on storage, networking, synchronisation among NSPs, and CPU dimensioning.

Finally, when an SLA request arrives, the service composition must first look into the database that contains all SLA offers to find the ones that could match the request. Above all, service composition must select the offer that could link the source to the destination of the SLA request.

As mentioned, this example illustrates one possible implementation to create offers for NSPs. ETICS will provide in further releases of the architecture also different approaches (with some recommendations) to match the different flavour of NSP's strategy to define their offers.

#### 5.2.2.1. Global SB Plane Scalability (push model) – Offer publishing

For the publish scenario (PUSH), the major problem comes from the scalability. As the number of NSP grows, the number of offers could become too large to be handled smoothly. Simulations and numbers extracted from real cases will more precisely determine, up to which number of autonomous systems, this scenario could be applied and under which conditions (frequency of service composition process call, acceptable time duration to obtain service composition results, offer validity time, etc.). For this PUSH scenario the subscription and notification service can be a possible scalable solution.

One possible approach to identify architectures and scalability properties for the service and business plane is to review similar solutions adopted in the World Wide Web, which is one of the main means for disseminating large amounts of information to users. The analogy that can be created is to consider the information about SLA as a pure contractual content (textual) that shall be exchanged (push model) between NSP and OTT and thus to apply consideration about web services related to content delivery architectures.

The ETICS push model for SLA can be organised in these phases:

1. Local storage of pushed offers at a local repository (NSP or facilitator)
2. Subscription of this repository (NSP or facilitator) by interested actors of alliance
3. Queries to the repository by subscribed actors[24]

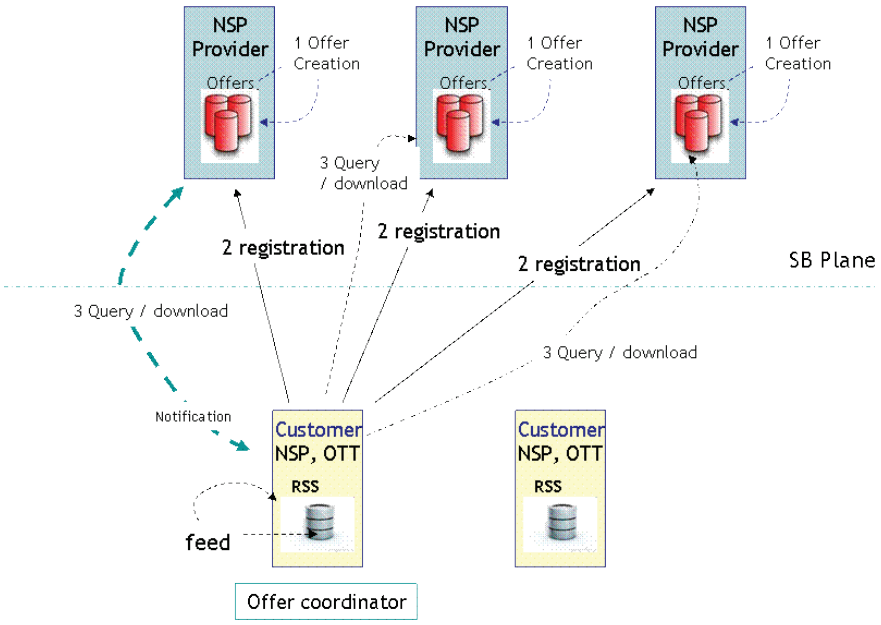


FIGURE 57: PUSH - SB PLANE OFFER DISSEMINATION

[24] The query to pre-compiled offers will be needed only towards the NSPs that are involved in the path computation by the coordinating entity (NSP, facilitator). Mechanisms like RSS [RSS] and WEB notification service [WNS] could be used for scalability reasons due to massive simultaneous queries to the repository, in order to have timely updated mirrors in the local domain of the coordinating entity.

This mechanism allows to have asynchronous mode between the local offer storing and the queries by subscribers such that the analysis of the scalability can be done for each of the three phases above where globally 1) will scale with the dimensioning of the storage system, 2) will scale up with the selected subscriber management system whilst 3) needs specific consideration for the 6 ETICS scenario to be done in next sections (see FIGURE 57).

The advantage is the fact that the synchronous listener on contents is not necessary at each NSP whilst all the query and download can be asynchronous to the SLA updates.

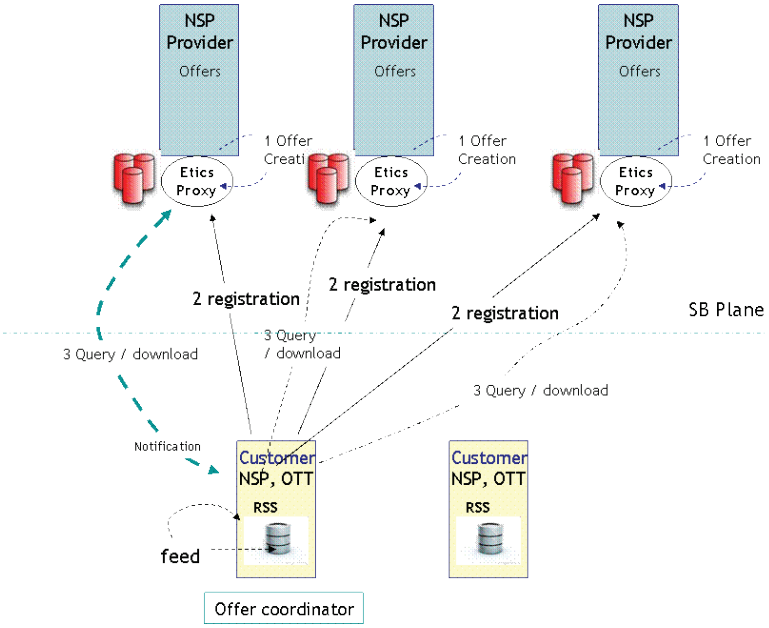


FIGURE 58: PUSH - SB PLANE OFFER DISSEMINATION INCL. INGRESS CONCEPT

The scalability can be ensured by introducing a proxy at the ingress of the NSP domain that will take care of the offer publishing and local storage (see FIGURE 58).

Although the collaboration and coordination among NSPs could be based on trusted private 'peer-to-peer' (i.e ERP, CRM and SCM) solutions[25] (that will have scalability issue driven by the vendors), we address the Web space that already has made progress on scalable architectures for collaboration among things and users having presentation and session layers consolidated interfaces (XML) and security solutions (HTTPS) to be proposed for B2B collaboration in ETICS.

In particular, as web-based services become more complex, the traditional Web model will become insufficient and thus all progress made on web for scalability can be adopted.

[25] ERP (Enterprise Resource Planning) utilises software applications to improve the performance of organizations' resource planning, management control and operational control. ERP software is a multi-module application software that integrates activities across functional departments. CRM (Customer Relationship Management) and SCM (Supply Chain Management) are two other categories of enterprise software that are widely implemented in corporations and non-profit organizations. While the primary goal of ERP is to improve and streamline internal business processes, CRM attempts to enhance the relationship with customers and SCM aims to facilitate the collaboration between the organization, its suppliers, the manufacturers, the distributors and the partners.

Although the literature and current implementation allows selecting a suitable scalable model (Push or Pull) to be adopted for ETICS here we try to provide an overview of *main key architectural aspects* that support large-scale, push-based and pull-based data delivery combined with customization in the section dedicated to all six ETICS scenarios.

### 5.2.3. On Demand Scenario (Pull Model)

As explained previously, in the On Demand Scenario, the NSPs are not publishing offers in advance. Instead, offers are computed according to the SLA request and published to others NSP (see FIGURE 59 below).

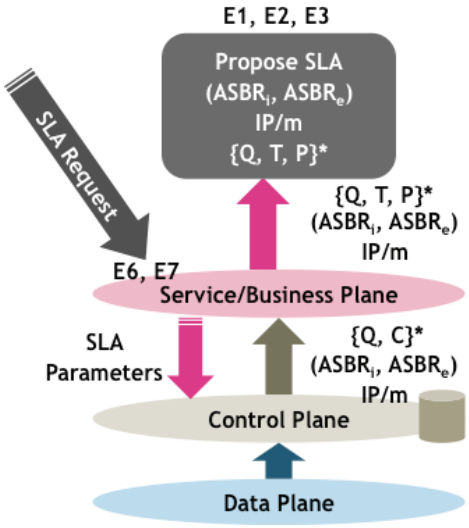


FIGURE 59: ON DEMAND SCENARIO FOR SLA COMPOSITION

In this scenario, an SLA request is proposed to the different NSPs which in turn, provide corresponding offers.

Compared to the publish scenario, the number of offers exchanged or notifications of offer updates between NSPs is limited and does not increase too fast when the number of NSP grows, but the service composition implies, at least, to propagate the SLA request between selected NSPs, with the aforementioned drawbacks:

- Imposing additional crank back mechanisms and more time required to obtain a result from the service composition process, or,
- Flooding the SLA request to all NSPs at the same time, this consumes more bandwidth and path computation power.

We consider that this approach is well designed when a 3rd Party like a broker (i.e. the ETICS consortium as a whole) centralizes the service composition process or if all the NSPs can benefit from an abstract global view to help at the selection of relevant NSPs to contact for the request. Through this fast overview on the different scenarios, both push and pull models have some advantages, but also weaknesses that prevent their application in all situations.

#### 5.2.3.1. Global SB Plane (SBP) Scalability (pull model)

Again all the Web based approaches made for the push model can be devised by reviewing typical pull based architectures and technologies for content delivery or messaging systems in the web space.

The ETICS pull model can be organised in these phases:

1. Customer NSP sends request for offer quotation to all NSP providers in the end-to-end path.
2. All interested NSP providers prepare the offer.
3. Interested providers send the respective offer.

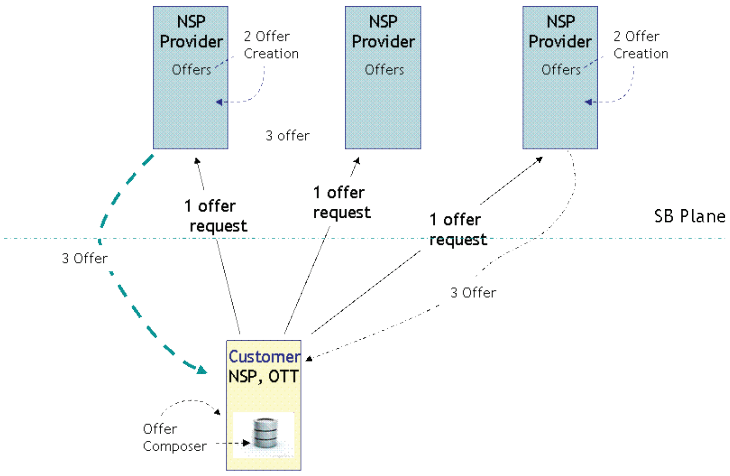


FIGURE 60: PULL OFFER DISSEMINATION

As said in the global introduction of Pull scenario there are few scalability issues to be considered for this modality and some consideration to be made about the performance and answering time from NSP provider when cascaded offers have to be compiled by the coordinator.

There are also security (D2.2 Req.: TR-SLA-COORD-DIST-03) issues to be considered for the Pull mode towards the NSP providers that may be involved in attachment. For this reason, the subscription mechanism is also preferred in this mode, to allow treatment of offer requests only to coordinators subscribed to the NSP. In this case, the Pull mode and queue mechanism will help to handle multiple pull offer requests. The only scalability consideration that has to be made by an NSP provider is due to the fact that it may receive several offer requests from different customers.

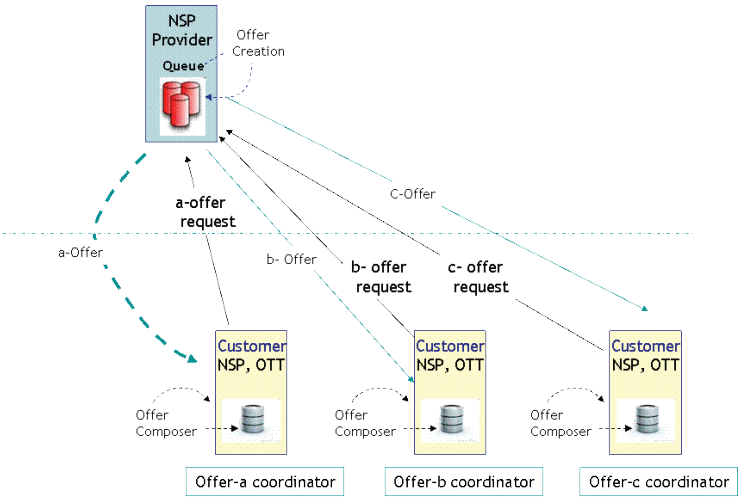


FIGURE 61: OFFER REQUEST QUEUE MANAGEMENT

The above FIGURE 61 highlights the need to handle an offer request queue by adding the subscription mechanism for the customer, while FIGURE 62 shows a possible scalable configuration of the NSP with the Proxy server handling both pull request queue and catalogue of pre-compiled offers.

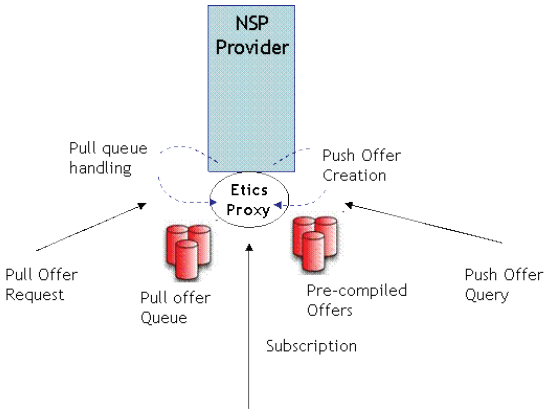


FIGURE 62: UNIVERSAL NSP CONFIGURATION FOR PULL AND PUSH MODE

### 5.2.4. SCALABILITY ANALYSIS OF THE SIX ETICS SCENARIOS

As explained previously, it is possible to provide further details for each of the scenarios on top of global consideration made for PUSH and PULL mode. In particular, the following tables try to give a solution to the requirement BR-GEN-12 (How the ETICS solution can scale from an initial implementation, involving only a few NSPs, where each pair of NSPs have a bilateral agreement, to a scale of several thousand independent actors).

| PUSH MODE | |
|---|---|
| **Fully centralized** by a **single centralized entity** "**facilitator**" for the whole community: offers (push) or service capabilities (pull)<br><br>Scalability analysis done for the push mode yields that the offer repository is in the facilitator and the need to have, for the pull request from various S, a handler of the queue.<br><br>The registration is not shown, the proxy is not shown at the facilitator | Fully Centralized – Push (4)<br> 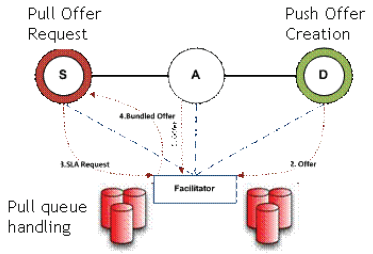 |
| **Centralized** by any NSP: instead of having a single facilitator, each NSP can be facilitator for some requests.<br><br>Scalability analysis done for the push mode yields that the offer repository is in the NSP coordinator and the queue handler as well.<br><br>The registration is not shown, the proxy is not shown at the facilitator | Per-NSP Centralized – Push (5)<br> 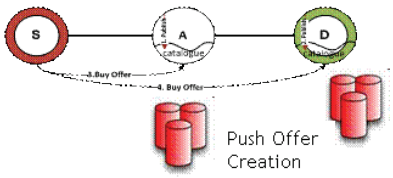 |
| **Distributed** through NSPs: in this case, each NSPs treats its part of the End-to-End demand and forwards the request with the remaining budget to following NSPs. The registration is not shown, the proxy is not shown at the facilitator | Distributed – Push model (6)<br> 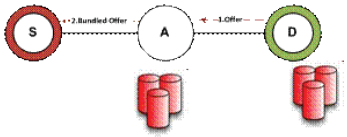 |

TABLE 1: OVERVIEW OF PUSH MODE ETICS SCENARIOS

| PULL MODE | |
|---|---|
| **Fully centralized** by a **single centralized entity** "**facilitator**" for the whole community: offers (push) or service capabilities (pull)<br><br>In this case there are no pre-computed offer and the facilitator only needs the queue handling for massive pull<br><br>The registration is not shown, the proxy is not shown at the facilitator | Fully Centralized – Pull model (1)<br> 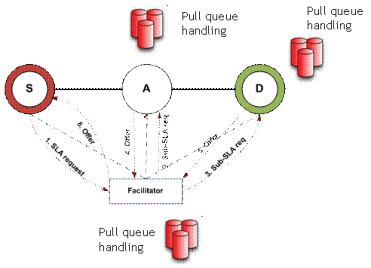 |

| | |
|---|---|
| **Centralized** by any NSP: instead of having a single facilitator, each NSP can be facilitator for some requests<br>The need to have queue handlers at each NSP<br>The registration is not shown, the proxy is not shown at the facilitator | Per-NSP Centralized – Pull(2)<br> 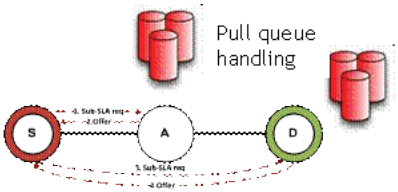 |
| **Distributed** through NSPs: in this case, each NSPs treats its part of the End-to-End demand and forwards the request with the remaining budget to following NSPs.<br>The registration is not shown, the proxy is not shown at the facilitator | Distributed – Pull models (3)<br> 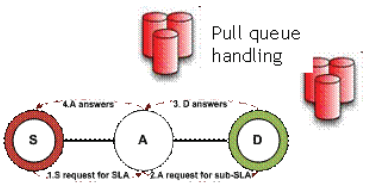 |

TABLE 2: OVERVIEW OF PULL MODE ETICS SCENARIOS

## 5.3. ECONOMIC FEEDBACK ON ETICS ARCHITECTURE EVOLUTION

Based on the feedback provided by [ETICS-D3.3], this section extracts the fundamental economic recommendations on ETICS architecture (as defined in [ETICS-D4.2]).

The deliverable [ETICS-D4.2] proposes six possible scenarios for ETICS solutions, which are organized around two axes:

- *Pull model* (per customer request, on demand) or *push* model (pre-packaged offers)
- Degree of centralization, i.e. *fully centralized* (independent and central coordination by a facilitator), *per-NSP centralized* (centralized coordination per NSP), or the *distributed* approach.

These two axes create the following set of *six combinations* as possible realisation variants for the ETICS architecture:

1. Push – Fully centralized
2. Pull – Fully centralized
3. Push – Distributed
4. Pull – Distributed
5. Push – Per NSP centralized
6. Pull – Per NSP centralized

In the following subsection we summarise [ETICS-D3.3]'s description of these architectural models. We then list the discussed business and economic criteria, while providing a tentative, qualitative evaluation of the six architecture alternatives. Open issues regarding these criteria are emphasised before we make a final summary of conclusions at the end of this section.

### 5.3.1. DETAILED SPECIFICATION OF THE COMPOSITION PHASE OF THE SCENARIOS

In D4.2 the authors have identified two different steps that must be followed in order for a service to be offered. These two steps consist of the *publishing* phase and *service composition* phase. In this section we describe as detailed as possible the composition phase of each of the six models. We assume that the publishing phase is already done and the NSPs in the community are aware of the network capabilities or the SLA offers of other NSPs. Also for the analysis below we assume that the buyer is an NSP that participates in the ETICS community and not an external end-customer.

The composition phase of the six models is shown in the figures below.

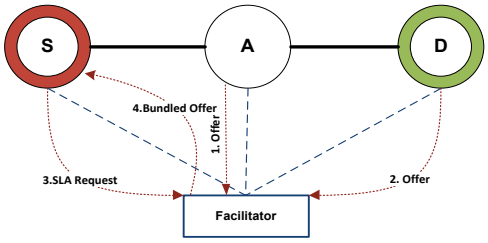


FIGURE 63 FULLY CENTRALIZED PUSH

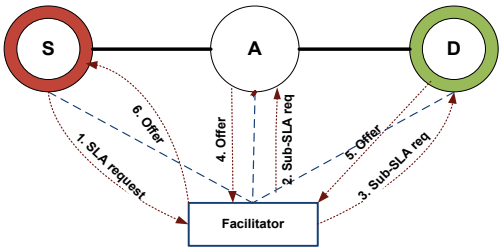


FIGURE 64 FULLY CENTRALIZED PULL

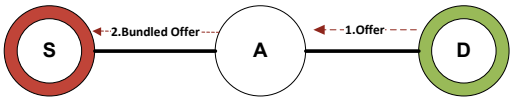


FIGURE 65 DISTRIBUTED PUSH

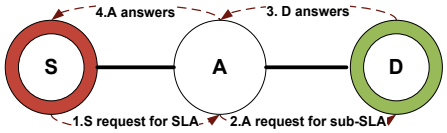


FIGURE 66 DISTRIBUTED PULL

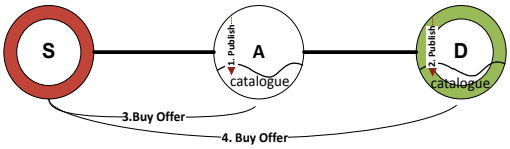


FIGURE 67 PER-NSP CENTRALIZED PUSH

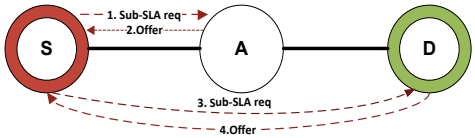


FIGURE 68 PER-NSP CENTRALIZED PULL

*Fully Centralized Push model*

In FIGURE 63 we show the steps that are followed in the composition phase of this model. Initially the NSPs that participate in the community agree to install a central entity to facilitate the composition phase (facilitator). All the NSPs in the alliance must acquaint the facilitator about their SLA offers (since it is a push model). Those SLA offers contain a logical point of interconnect (PoI), on-net destination(s), QoS characteristics, and expiration time of the offer and the price. The buyer (S) communicates with the facilitator and reveals only the PoI, the destination, and the QoS characteristics that he wishes to have for the specific connectivity (SLA Request). After receiving a proposed ASQ good from the facilitator (Bundled Offer), he can accept or reject it. The facilitator is aware of the various offers that the NSPs have already provided. It is also aware of what the buyer is willing to acquire. Thus based on this knowledge it combines offers of those NSPs that form a chain from the source NSP to the destination NSP(s).

*Fully Centralized Pull model*

The facilitator in this model (FIGURE 64) is aware of the network capabilities of each NSP in the community due to the publishing phase. Based on this knowledge, the facilitator computes one or many NSP chains that could potentially handle the buyer request for PoI, destination and QoS characteristics. It then sends specific requests to the NSPs in those chains. After receiving the SLA offers, it combines them and chooses one to return to the buyer.

For both fully centralized models, in order for the facilitator to be able to combine offers, the objective under which it is functioning has to be determined by the NSPs participating in the community and also to be known to all NSPs participating in the community. Nevertheless, its main objective is to provide the buyer with offers that are compliant with his QoS requirements and also are likely to meet his price.

*Distributed Push model*

Depending on the publish scenario that is used, each NSP either knows the SLA offers of all the NSPs in the community or just a subset published to it. These offers contain PoI, on-net destination(s), the QoS characteristics, the expiration time of the offer and a price. If an NSP knows all the SLA offers available in the community then, in case of becoming a buyer, he can combine offers and create an NSP chain. If the propagation of the offers is made only to a subset of NSPs such as the direct neighbours, a way of combining offers has to be determined. This means that each NSP according to his own strategy propagates combined offers to other NSPs in order for the information to be diffused to each and every participant. In any case, the buyer will buy the whole service from its neighbour NSP, who will negotiate with his own neighbour etc. Thus even if the buyer may know about the SLA offers from all NSPs in the community he cannot buy from them directly but has to rely on his neighbours to combine specific offers (FIGURE 65).

*Distributed Pull model*

The buyer requests connectivity from its neighbour(s) in order to reach a destination (FIGURE 66). Each NSP accepts (or rejects) this SLA Request. In case of accepting, he extracts the part of the requesting SLA that corresponds to his network capabilities and adds a price. He propagates constraints on price and remaining network capabilities to his neighbour(s).

*Per-NSP Centralized Push model*

Each NSP creates SLA offers that contain the on-net destination(s), the QoS characteristics, the expiration time of the offer and a price. The buyer, who knows about those offers due to the publishing phase, combines offers and creates a bundled one that fulfils his needs. As opposed to the distributed model, here the buyer buys the SLA offers directly from the NSPs even if they are not directly connected to him (FIGURE 67).

*Per-NSP Centralized Pull model*

The buyer asks each NSP in the chain for an SLA offer (which is a part of the whole ASQ good that the buyer wants). In this model (FIGURE 68) the buyer communicates separately with each NSP in the chain. The apportioning of price and/or QoS is done according to the central NSP's (buyer) own strategy who leads the composition. In certain cases, opportunities for service provision can be missed, although a different central NSP could have been successful in meeting the buyer objectives.

Responsibility of the SLA commitment

In the various service composition models, different entities are responsible for the end-to-end SLA commitment: In the fully centralized models, the facilitator is the only entity responsible for the end-to-end SLA commitment, because the contract is established between it and the buyer NSP. In case of a problem, the facilitator may interrogate each NSP that is involved in the contract to determine the faulty network. In the Distributed models the first contract is established between the buyer and the seller. Thus the latter has the responsibility of the end-to-end SLA commitment. Since each NSP in the chain subcontracts his neighbour in order to form buyer-seller pairs until the destination is reached, the seller NSPs is always responsible. In Per-NSP Centralized models the buyer is responsible for the contracts it establishes with each other NSP. The end-to-end SLA commitment is guaranteed only by the buyer as other NSPs are not aware of the usage of their respective offers in the service composition. In case of failure the buyer is responsible to determine the faulty domain.

### 5.3.2. OPEN ISSUES IN DESCRIPTION OF THE MODELS

Some open points have limited the economic and business analysis of [ETICS-D3.3] regarding the suitability of [ETICS-D4.2]'s architectural models.

[ETICS-D3.3] has highlighted the imprecise separation of terms like network capabilities and SLA offers. The analysis of [ETICS-D3.3], hence, had to rely on the following hypothesises:

- *Network capabilities do not seem to make sense without defining the basic connectivity endpoints (source and destinations), if any form of pre-selection of paths is required to determine which NSPs to include in the final offer composition. The source in this context represents the logical Point of Interconnection (PoI).*
- *An "SLA offer" only requires the acceptance from the buyer in order to be put into operation; an offer is not renegotiated. Hence, it must contain all details required for a final agreement – especially unambiguous and un-disputable pricing information. From a business point of view, scenarios mainly differ regarding which details are communicated with whom and when in the process of creating the composed SLA offer, in particular regarding QoS information.*

Based on this recommendation of [Del3.3], the description of network capabilities and SLA offers has been enhanced in Section 4 of the present deliverable on the basis of contributions from WP4 and WP5.

The rest of the open issues addressed in this section will however only be addressed in the next ETICS architecture deliverable D4.4.

[ETICS-D3.3] has assumed that the publishing phase has already been concluded. In this recommendation, the following four models of publishing alternatives had been discussed [ETICS-D3.3]:

1. *Publish only to direct neighbours. This option protects confidentiality, since it does not allow a global vision of what SLA offers or network capabilities are available inside the alliance. Due to its nature, this is only pertinent in the distributed scenarios.*
2. *Publish only to the facilitator. This option is only available in the Fully Centralized scenarios.*

3. *Publish to all NSPs within the alliance. In this case, the architecture must support a flooding mechanism that can overcome potentially non-cooperative NSPs.*

4. *Publish to a subset of NSPs governed by policy rules. This is a generalisation of the three options above and might need complicated mechanisms from a technical point of view.*

*Fully Centralized models*

In the Fully Centralized models we described the composition phase assuming that the way the facilitator chooses NSPs to form a path is a black box operation. This process of the facilitator affects the model and thus it has to be clarified. There are three possible ways for the facilitator to compute a path:

1. Choose NSPs according to alliance criteria: The facilitator attempts to maximize the community's welfare by taking into account criteria such as the total number of satisfied customers or the load-balancing of the network.

2. Choose NSPs according to buyer's criteria: This is source routing. One of the advantages of source routing is that the source (here the buyer NSP) chooses a route according to its own objectives. By definition QoS requirements are constraints and price, number of buyers and load-balancing are all potential component of an NSP's objective function.

3. Provide all routes and let the buyer decide. This is very similar to operating the facilitator according to buyer's criteria (option 2), but opens up for more advanced strategic behaviour from the buyers as they are not limited by facilitator functionality.

*Fully Centralized Pull model*

In the description of the Fully Centralized Pull model, we assumed that somehow the facilitator requests parts of the SLA from the NSPs that participate in a chain, in order to compose a bundled offer. In fact the splitting of the SLA request of the buyer can be done in at least three ways:

1. The facilitator requests the first NSP in the NSP chain to execute a Backward Recursive Path Computation to determine the ASQ offers. The AS chain is given as input parameter.

2. The facilitator provides the first NSP in the chain with the global SLA request (the one coming from the buyer). When it receives the offer, it forwards it to the second NSP and so on. The offer contains the remaining QoS budget and the Point of Interconnect chosen by the current NSP, where the next one must start.

3. The facilitator requests each NSP based on the fixed network capabilities for each NSP, the QoS budget and the Point of Interconnect (PoI).

*Per-NSP Centralized Pull*

In a similar way as above, the central NSP in the Per-NSP Centralized model has to determine the way of requesting different parts of an SLA request from the NSPs in a chain:

1. By using the PCE BRPC algorithm (in a distributed way) with the NSP previously computed (including the Point of Interconnect).

2. By asking each NSP in the chain to provide their part of the ASQ good sequentially by passing the results of NSP n to NSP n+1.

3. By asking each NSP in the chain to provide their part of the ASQ good by imposing QoS budget and the Point of Interconnect.

*Distributed Push*

In the Distributed Push model, the way the publishing phase is done, affects the information that is distributed in the community. In case of using a subset or only the direct neighbours for publishing, the NSPs in the community will not be aware of all SLA offers that are available. Thus in this case NSPs must combine their offers with their neighbours' ones in order for the NSPs in the community to have "global" reachability. The way, how such combinations could be done may have to be described.

### 5.3.3. Recommendation of Architectural Models

In this subsection, economic and business recommendations of D3.3 [ETICS-D3.3] are extracted, which provide indications when specific architectural models seem to be preferable over others. Based on pure economic analysis the following three recommendations declare the Fully Centralised Push option to be most beneficial:

| | |
|---|---|
| ***EC-REC-1*** | *From a pure economic perspective per-NSP centralised architectures may be disadvantageous. Large NSPs might be able to use their large customer base, network footprint and service offerings to bundle ASQ goods with best effort products or other goods and price them strategically to gain market advantages. From the same economic perspective, the fully centralized push option is preferred.* |
| ***EC-REC-2*** | *In a mature market, push is preferred over pull. In addition, distributed architectures may be of interest* |
| ***EC-REC-3*** | *The costs in building and maintaining a centralized facilitator, including management of rules and security, is not believed to be high compared to the advantages above.* |

This argument is moreover supported by low cost estimations for using a centralised facilitator unit. On the other hand, [ETICS-D3.3] has drawn our attention to the difficulty of realising such architectural models in practice:

| | |
|---|---|
| **BI-REC-1** | *Fully centralised architectures need very clear business incentives in order to make it attractive for operators to handover the composition process to a third party.* |
| **BI-REC-2** | *We see the per-NSP centralized pull as the most probable first step for the implementation of the ETICS architecture.* |
| **BI-REC-3** | ***Therefore ETICS should consider the per-NSP centralized pull option as a possible decisive first step, as well as hybrid models which may be possible options in the longer term*** |

As a consequence, we can assume that the per-NSP centralised pull model will most likely serve as starting point bootstrapping the market. Hence, it may be seen as fundamental architectural starting point as well. As stated above, push models may become more attractive when the market matures. As an overall result, D3.3 has neither been able to extract a dominating architectural model, nor been able to completely eliminate any models, but has rather highlighted their relative strengths and weaknesses.

### 5.3.4. Coexistence of Architectural Models

According to [ETICS-D3.3], the heterogeneity of the ETICS community and its offered services in practice may be captured along two dimensions. On the one hand, more than one architectural model may be applied by NSPs in parallel, i.e. some may cooperate in a distributed manner as known from today's Internet and some may seek for varying degrees of centralisation. On the other hand, independent herds of cooperation or centralisation may be formed.

| | |
|---|---|
| **BI-REC-4** | *If several NSPs choose to follow the per-NSP approach for specific customers and services, we should be aware that this may create pockets of incompatible ASQ inter-carrier solutions. Large quantities of such pockets will not be efficient in providing a significant quality-assured service offering. In a second step it may be possible to see some brokers enter the market and contact carrier communities, asking for offers on a pull basis.* |
| **BI-REC-5** | *ETICS should consider the case where similar trusted NSPs choose a centralized architecture while others form a different community using distributed solutions with minimal information sharing.* |

## 5.4. PERFORMANCE ANALYSIS

According to the work done towards [ETICS-D5.4], the distribution of the simulations cases follow clear patterns, described in following paragraphs.

The specificity of ETICS is to join technical and economic points of view. The goal is to provide a way to obtain a path in the network crossing different carriers, economically negotiated, insuring QoS and network services guarantees. This is the reason why it has been decided that the simulations focus on three consecutive aspects of the related process, which has led to the three themes proposed in [ETICS-M5.3] section 1, based on the [ETICS-D4.2] architecture description:

I. Technical-economic negotiations and strategies (prices, QoS level, SLA, …)

II. Path selection

III. Path establishment, control and use

In order to supply feedback to the present document, some partial results have been extracted from the individual simulations. These recommendations extracted from partial simulations are not definitive, as the simulations will be developed also during the last 12 months of the project after [ETICS-D4.3] completion.

Up to date, individual results and conclusions have been extracted in path selection and path control and specific simulations related with survivability, which is reflected in following subsections.

### 5.4.1. PATH SELECTION

#### 5.4.1.1. Efficiency and cooperation in the routing of flows: Simulations study and impact assessment

It has been studied both by means of a theoretical model and simulations whether simple policies in the way user session requests are handled have a positive or negative impact on the overall ETICS system and the user due to the impact on the call blocking probability, as well as the individual NSP monetary rewards. The user strategies Nash equilibrium properties are analysed and the impact of deviation is quantified.

In particular, our simulations consider a *fully connected*, *symmetrically loaded* network of *N* NSPs. The routing strategy *s_i* is defined as trying up to *i* (randomly chosen) two-hop alternative path routes, with *k* denoting the maximum number of alternative two-hop paths when a direct one-hop path is not available. The *least loaded* routing policy mentioned henceforth is a variation of the *s_k* policy where a flow, if not routed directly, is routed via the least loaded alternative two-hop path. We briefly state some major simulation scenarios run and the respective results in the remainder of this subsection.

Sample simulation scenarios are specified as follows:

- **Scenario A:** All ETICS NSPs adopt the myopic (direct) routing policy.
- **Scenario B:** All ETICS NSPs adopt the *s_k* routing policy.
- **Scenario A':** All but one ETICS NSPs adopt the myopic (direct) routing policy, one the *s_k* policy.
- **Scenario B':** All but one ETICS NSPs adopt the *s_k* routing strategy, one the myopic policy.
- **Scenario C:** All ETICS NSPs adopt the least loaded routing policy.
- **Scenario C':** All but one ETICS NSPs adopt the least loaded routing policy, one the myopic (direct) routing policy.
- **Scenario C'':** All but one ETICS NSPs adopt the myopic (direct) routing policy, one the least loaded routing policy.

The definition of these scenarios allows us to cross-compare the performance of each NSP when adopting a common routing policy or deviating from it, hence quantify the Nash equilibrium properties (or cost of deviation from equilibrium) and assess the overall system performance and end user satisfaction by means of computing the blocking probability of the system.

It is shown that indeed *s_k* comprises a Nash equilibrium strategy over the strategy set *{s_k, myopic}*. This is also quantified by the simulations which indicate that the reward of the (deviating) NSP that adopts the *myopic* policy under Scenario B' is less than that under Scenario B. Similarly, it is shown that the *myopic* policy does not comprise a Nash equilibrium strategy over the strategy set *{s_2, myopic}*. This is also quantified by the simulations which indicate that the reward of the NSP that adopts the *myopic* policy under Scenario A' is greater than that under Scenario A. The blocking probability is also assessed and it is

depicted that the blocking probability under the Scenario A is less than the Scenario B. This results in higher benefit for the users who receive better service from the ETICS networks infrastructure.

By working with the same methodology, useful insight is provided with respect to the impact of each routing strategy. In most simulation runs, the *s_k* routing policy and the least loaded routing policy perform really close, despite the lack of sophistication of the former strategy as opposed to the latter, which also needs more information regarding the state of the network.

Also it is concluded that

a) given that the network is not under-utilised, and b) over a set of routing policies which require different amounts of information regarding the overall network state, selfishness and greedy policies mitigate both system performance and rewards, regardless of the amount of information shared or the sophistication of the routing policies.

On the other hand, agreeing on common rules, which do not necessarily comprise Nash equilibrium strategies, such as prioritization of high-value flows and accepting low-value flows only when a capacity threshold is available over an ASQ path, is beneficial for the overall system performance and blocking probability. This provides useful guidelines for the way session flows can be efficiently routed on top of the ETICS ASQ infrastructure. And from these guidelines the following recommendations:

**SUMMARY OF *SIMULATION 1* RESULTS**

- If the ETICS ASQ infrastructure is saturated, then it is important to prioritise single-hop high-value flows as opposed to low-value or multi-hop flows. Doing the latter, though it will be myopically beneficial for the NSP serving the customer, will reduce the network efficiency (the system blocking probability deteriorates) and thus the number of users served. Therefore, a common rule for CAC on top of ASQ goods is needed where all NSPs agree to accept low-value flows only when there is enough free capacity (above a threshold).
- Implementing a trunk reservation policy in the centralized architecture would be highly beneficial in terms of system performance. The central control and full information over the resources available on the network and the market demand exhibited allows for a more efficient allocation of resources to the competing user flows, improving the system blocking probability that is experienced by all the ETICS users.
- Overall, *it is highly recommended that the deployment of the ETICS architecture is combined with an agreement on system common rules regarding the admittance of flows and routing*. This recommendation applies to all 6 different options of the ETICS architecture.

#### 5.4.1.2. IC coordination models and information issues: Simulations study and comparison of different architecture models

In [ETICS-D5.4] we have provided a detailed simulation study of the IC coordination and information issues and the way these impact the ETICS system performance. In particular, for each of the architecture options we have investigated the information available in each NSP that participates in the community. For details of this work and simulation results we refer to [ETICS-D5.4] (the modelling has also been used in WP3).

For completeness reasons, we would like to note that the information set that each participant will have available in each model depends on:

- the NSPs to which a participant may send information;
- the smallest and the largest piece of information that can be shared;
- the information that is shared from other NSPs and propagated to others;
- the strategy of the buyer (revealing his willingness to pay or not);
- the pricing strategy of the central entity (if any); and
- the propagation of the offers.

Due to those dependencies, various possible sub-models may result under each of the models that we previously mentioned. The analysis of these models and their simulation comprises on-going work. So far, we have focused on the distributed pull model.

**SUMMARY OF *SIMULATION 2* RESULTS**

- The simulation of the distributed pull model indicates that selfishness has a high toll on the end user satisfaction, since the **probability that the service is offered** when the first NSP acts selfishly is less than that in the **collaborative model**. Therefore, the ETICS architecture and in particular the ETICS community under the distributed pull architecture option should enforce rules for collaboration and forbid/punish aggressive selfish pricing strategies that mitigate the ETICS system's performance and the interest of the ETICS community as a whole.

#### 5.4.1.3. IC paths under QoS and price constraints: Path Selection

Algorithms have been tested to build IC paths under QoS and price constraints. The algorithms are designed to be used in the ETICS Network at the control Plane.

In the second case of our simulations, algorithms are used in the pull model (centralized or distributed) when the network capabilities exchanged between NSPs do not allow for knowing the QoS budget per domain (e.g. if network capabilities only allows to discover the inter domain topology). In this case, the composing entity (facilitator in the centralized model, or the ingress NSP in the distributed model) needs to ask each possible NSP for a specific offer with respect to the customer demand. We propose and demonstrate the possibility to extend the PCE architecture by replacing the standard path computation method (BRPC) by a new protocol and a related algorithm allowing the simultaneous exploration of multiple inter-domain paths in order to select best NSP offers for the End-to-End request. In this case, each domain PCE computes a path between its border nodes and forward one or several paths to following PCEs. The tested topology includes 5 domains and hundreds of nodes per domain, and requests include 3 to 5 QoS constraints. Results show that:

**SUMMARY OF *SIMULATION 3* RESULTS**

- Obviously exploring multiple paths (instead of on AS chain in the standard BRPC approach) allows accepting more requests satisfying the end-to-end QoS constraints, even if only the best path is kept at each decision step and forward to next NSP.

**SUMMARY OF *SIMULATION 4* RESULTS**

- As a trade off between the quality of the answer, the overhead in terms of message and the execution time, storing a single end-to-end path (among the various multiple paths explored) seems the right choice, but also that storing more than two paths is not required (i.e. shows no significant benefit).

**SUMMARY OF *SIMULATION 5* RESULTS**

- This execution time is kept quite reasonable for five crossed autonomous systems. Intrinsically, the method is faster than having multiple BRPC sessions to explore the same AS path diversity. In the worst case, when the AS paths are not completely disjoint from the source AS to the destination AS, the solution is similar to BRPC in terms of execution time. However, in most cases, when one or several ASes are part of different AS chains, our solution allows to compute paths only once for ASes in common, while computations are done once per chain with BRPC. Our solution scales approximately linearly with the number of explored ASes, while it scales more exponentially with multiple BRPC sessions in parallel, each chain being independent.

### 5.4.2. PATH ESTABLISHMENT, CONTROL AND USAGE

#### 5.4.2.1. SLA Validation

Following the criteria in the documents generated in T5.2 [ETICS-M5.3] Section 5.4, [ETICS-M5.5] slide 26, and architecture in [ETICS-D4.2], an environment for test and simulation cases of SLA violation has been established. The purpose is to obtain results to validate possible scenarios in the control of the path establishment. According to the previous assertion, an environment for test to validate a specific SLA agreement with QoS parameters into an IP network has been setup through testing and evaluation traffic of different parameters between two networks.

A client generates a desired amount of data stream that floods the network to measure whether a specific SLA meets the conditions specified according to SLA parameters. This procedure gave us the amount of packet loss, jitter, etc, specifying whether the QoS is satisfied for a previously established SLA.

In the test we could obtain some conclusions about the behaviour of the interconnection between two or more NSPs that share interconnection traffic. After the previous testing the following recommendation has been extracted:

**SUMMARY OF *SIMULATION 6* RESULTS**

- It is possible in the ETICS community that some NSPs suffer changes in the state of their network, producing from time to time problems in the fulfilment of SLA. Therefore, SLA violation control methods in ETICS architecture are necessary.
- Methods are needed to find the best NSPs (in parameters, QoS, Jitter, Delays, etc.) which is to be included in the creation of a product.

# 6. CONCLUSIONS AND OUTLOOK

This document presents the current status of the ETICS architecture as a revision of the architecture which has been presented in deliverable *D4.2 – ETICS Architecture and Functional Entities High Level Design [ETICS-D4.2]*. The results of this work are based on the continuous work stream in the ETICS architecture work package (WP4) as well as the inputs provided by work packages which address economic issues (WP3) and which are in charge of designing detailed technical specifications of the ETICS system (WP5).

First, we have identified the main constraints and assumptions on internetworks, followed by the presentation of the ETICS overlay model along with the most important high-level options for cooperation models among ETICS operators. After introducing the main types of ETICS services, the central part of D4.3 presents the ETICS reference architecture along with the related service deployment scenarios. In contrast to [ETICS-D4.2], this document has strongly focused on technically modelling architectural details with the help of Unified Modelling Language (UML) diagrams. The investigated deployment scenarios have demonstrated a series of implications on the implementation, and have further clarified various details such as the distributed nature of the centralised entity's composition functions. Moreover, our work has indicated the requirement of solid and independent monitoring mechanisms, which are capable of measuring and verifying whether an SLA has been satisfied.

Finally, the presentation of the ETICS system is followed by an analysis of the architecture's scalability, performance, and economic purposefulness, which has resulted in valuable architectural recommendations to be addressed in future revisions.

While the ETICS architecture – as presented in this report – has reached a high level of maturity, serving as the basis for the detailed specification of ETICS components and their implementation in the ETICS testbed, the architecture specification work has not yet come to an end. In a follow-up to this deliverable, *D4.4 - Final ETICS architecture and functional entities high level design* will incorporate all changes to the ETICS architecture which will be based both on the continued work within the architecture work package itself as well as on inputs by the related ETICS activities, and it will be delivered in Month 34 (M34) of the ETICS Project.

# 7. REFERENCES


[ACTRICE-D1.3] ACTRICE Deliverable D1.3, *Les alliances comme mode d'organisation économique pour la fourniture de services en inter-domaine*, 2007.

[BGPMEA] Huston, G., *BGP Routing Table Analysis Report*, Web site: http://bgp.potaroo.net/.

[BoLe05] Boucadair, M. and Levis, P., *De nouvelles perspectives pour la gestion de la Qualité de Service inter-domaine à grande échelle*, In Proceedings of the GRES'05 Colloquium, Luchon, France, March 2005.

[CL02] Contractor F. J., Lorange P., *The growth of alliances in the knowledge-based economy*, International Business Review, No 11., 2002.

[ETICS-D2.1] Economics and Technologies for Inter-Carrier Services (ETICS), INFSO-ICT-248567, Deliverable D2.1 – *Current business models and services; scenarios for the future; high-level requirements – How can the future Internet look like?*, May 2010.

[ETICS-D2.2] Economics and Technologies for Inter-Carrier Services (ETICS), INFSO-ICT-248567, *Deliverable D2.2 – Business and technical requirements for future network architectures*, January 2011.

[ETICS-D3.2] Economics and Technologies for Inter-Carrier Services (ETICS), INFSO-ICT-248567, *Deliverable D3.2 – Potential business models analysis and requirements*, January 2011.

[ETICS-D3.3] Economics and Technologies for Inter-Carrier Services (ETICS), INFSO-ICT-248567, *Deliverable D3.3 – Financial/Economic Dynamic Analysis*, December 2011.

[ETICS-D4.1] Economics and Technologies for Inter-Carrier Services (ETICS), INFSO-ICT-248567, *Deliverable D4.1 – End-to-End service specification template*, November 2010.

[ETICS-D4.2] Economics and Technologies for Inter-Carrier Services (ETICS), INFSO-ICT-248567, *Deliverable D4.2 – ETICS architecture and functional entities high level design*, June 2011.

[ETICS-D4.4] Economics and Technologies for Inter-Carrier Services (ETICS), INFSO-ICT-248567, *Deliverable D4.2 – Final ETICS architecture and functional entities high level design*, expected December 2011.

[ETICS-D5.2] Economics and Technologies for Inter-Carrier Services (ETICS), INFSO-ICT-248567, *Deliverable D5.2 – ETICS Draft Detailed specification of the inter-carrier service delivery system*, December 2011.

[ETICS-D5.3] Economics and Technologies for Inter-Carrier Services (ETICS), INFSO-ICT-248567, *Deliverable D5.3 – First release of selected components for the Inter-Carrier WP5*, 2011.

[ETICS-D5.4] Economics and Technologies for Inter-Carrier Services (ETICS), INFSO-ICT-248567, *Deliverable D5.4 – Simulative Assessment on ETICS intercarrier Service Delivery solution*, 2011.

[ETICS-D6.1] Economics and Technologies for Inter-Carrier Services (ETICS), INFSO-ICT-248567, *Deliverable D6.1 – ETICS Testbed Specification and Implementation*, May 2011.

[ETICS-M5.2] Economics and Technologies for Inter-Carrier Services (ETICS), INFSO-ICT-248567, *Milestone M5.2 – Selection of simulation environments*, 2010.

[ETICS-M5.3] Economics and Technologies for Inter-Carrier Services (ETICS), INFSO-ICT-248567, *Milestone M5.3 – Detailed specification of the algorithms to be deployed in the solution*,

[ETICS-M5.5] Economics and Technologies for Inter-Carrier Services (ETICS), INFSO-ICT-248567, *Milestone M5.5 – Simulators ready*, 2011.

[ETSI-1] European Telecommunications Standards Institute (ETSI) Telecommunications and Internet converged Services and Protocols for Advanced Networking (TISPAN), *Resource and Admission Control Subsystem (RACS): Functional Architecture*, Technical Report ETSI ES 282 003 v3.4.2, 2010.

[ETSI-2] European Telecommunications Standards Institute (ETSI) Telecommunications and *Internet converged Services and Protocols for Advanced Networking (TISPAN): Network attachment subsystem (NASS)*, ETSI ES 282 004 v3.4.1, March 2010.

[GaRe01] Gao, L., Rexford, J., *Stable Internet routing without global coordination*, IEEE/ACM Trans. Networking Vol. 9, No. 6, December 2001.

[GrSh02] Griffin, T., Shepherd, F. B., Wilfong, G., *The stable paths problem and interdomain routing*, IEEE/ACM Trans. Networking Vol. 10, No. 1, 2002.

[IETF-DR-1] Internet Engineering Task Force (IETF), Network Working Group, *Advertisement of Multiple Paths in BGP – Internet Draft*, Work in Progress, ed. Walton, D., Chen, E., Retana, A., September 2011.

[IETF-DR-2] Internet Engineering Task Force (IETF), *Network Best Practices for Advertisement of Multiple Paths in BGP – Internet Draft*, Work in Progress, ed. Uttaro, J., Van den Schrieck, V., Francois, P., Fragassi, R., Simpson, A., Mohapatra, P., October 2010.

[IETF-DR-3] Internet Engineering Task Force (IETF), Network Working Group, *North-Bound Distribution of Link-State and TE Information using BGP – Internet Draft*, Work in Progress, ed. Gredler, H., Medved, J., Farrel, A., Previdi, S., September 2011.

[IETF-DR-4] Internet Engineering Task Force (IETF) Draft Y.1731-07, MPLS-TP OAM based on Y.1731, ed. Busi, I., and van Helvoort, H., 2011.

[IMS] 3rd Generation Partnership Project, IP Multimedia Subsystem (IMS), [Online]. Available: http://www.3gpp.org/article/ims

[IPSPH-R1] IPsphere 1.0, IPsphere Framework Technical Specification (Release 1). June 2007.

[IPSPH-TR158] [IPsph-TR158] TR158, *IPsphere Framework, General Requirements and Technical Architecture*, Release 2.0.

[ITU2012] International Telecommunication Union (ITU), *Recommendation ITU-T Y.2012 (04/2010) – Functional requirements and architecture of next generation networks*, 2010.

[JaCo11] JAIN Community, February 2011. [Online]. Available: HTTP://JAVA.SUN.COM/PRODUCTS/JAIN

[LISP] Internet Engineering Task Force (IETF), *Locator/ID Separation Protocol (LISP)*, IETF Internet Draft (Work in Progress), ed. Lewis, D., Fuller, V., Farinacci, D., Meyer, D., January 2012.

[M.1400] International Telecommunication Union (ITU), *Recommendation ITU-T M.1400 – Designations for interconnections among operators' networks*, July 2006.

[M.3100] International Telecommunication Union (ITU), *Recommendation ITU-T M.3100.– Generic network information model*, April 2005.

[NeHa] Neginhal, M., Harfoush, K., and Perros, H., *Measuring Bandwidth Signatures of Network Paths*, In Proceedings of the 6th International IFIP-TC6 Conference on Ad Hoc and Sensor Networks, Wireless Networks, Next Generation Internet (NETWORKING'07), Springer, 2007.

[OASIS] *Advanced Message Queuing Protocol (AMQP)* [Online]. Available: http://www.amqp.org/

[OMA] Open Mobile Alliance, *Next Generation Service Interfaces Architecture* [Online], Available: http://member.openmobilealliance.org/ftp/ Public_documents/ARCH/Permanent_documents/OMA-AD-NGSI-V1_0-20100401-D.zip

[OnAp] *GSM World One API*, March 2011, [Online]. Available: http://www.gsmworld.com/oneapi/

[Or11] Oracle, *JAIN and Java in Communications*, February 2011, [Online]. Available: http://java.sun.com/products/jain/reference/docs/Jain_and_Java_in_Communications-1_0.pdf

[Or11-2] Oracle, *JAIN and Open Networks*, February 2011, [Online]. Available: http://java.sun.com/products/jain/JainAndOpenNetworks01.pdf

[PaX] *Parlay X*, March 2011. [Online]. Available: http://www.parlayx.com/.

[RFC2679] Internet Engineering Task Force (IETF), Advanced Network & Services, *RFC 2679 – A One-way Delay Metric for IPPM*, ed. Almes, G., Kalidinidi, S., and Zekauskas, M.,, September 1999.

[RFC2680] Internet Engineering Task Force (IETF), Advanced Network & Services, *RFC 2680 – A One-way Packet Loss Metric for IPPM*, ed. Almes, G., Kalidinidi, S., and Zekauskas, M., September 1999.

[RFC4271] Internet Engineering Task Force (IETF), Network Working Group, *RFC 4271 – A Border Gateway Protocol 4 (BGP-4)*, ed. Rekhter, Y., Li, T., and Hares, S., January 2006.

[RFC4656] Internet Engineering Task Force (IETF), Network Working Group, *RFC 4656 – A One-way Active Measurement Protocol (OWAMP)*, ed. Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and Zekauskas, M., September 2006.

[RFC4655] Internet Engineering Task Force (IETF), X, *RFC 4655 – A Path Computation Element (PCE)-Based Architecture*, Farrel, A., Vasseur, J.-P., Ash, J., August 2006.

[RFC5101] Internet Engineering Task Force (IETF), Network Working Group, *RFC 5101 – Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information*, ed. Claise, B, January 2008.

[RFC5136] Internet Engineering Task Force (IETF), Network Working Group, *RFC 5136 – Defining Network Capacity*, ed. Chimento, P., Ishac, J., February 2008.

[RFC5160] Internet Engineering Task Force (IETF), Network Working Group, *RFC 5160 – Considerations of Provider-to-Provider Agreements for Internet-Scale Quality of Service (QoS)*, ed. Levis, P., and Boucadair, M., 2008.

[RFC5441] Internet Engineering Task Force (IETF), Network Working Group, *RFC 5441 – A Backward-Recursive PCE-Based Computation (BRPC) Procedure to Compute Shortest Constrained Inter-Domain Traffic Engineering Label Switched Paths*, ed. JP. Vasseur, R. Zhang, N. Bitar and J.L. Le Roux, 2009

[RSS] Wikipedia, *Web Feed*, [Online]. Available: http://en.wikipedia.org/wiki/Web_feed

[W3C11] World Wide Web Consortium (W3C), *Simple Object Access Protocol*, February 2011, [Online]. Available: http://www.w3.org/TR/soap/

[WAC] *Wholesale Application Community (WAC) Platform*, March 2011, [Online]. Available: http://public.wholesaleappcommunity.com/

[WNS]: The 52° North, *Web Notification Service (WNS)*, [Online]. Available: http://52north.org/communities/sensorweb/wns/1.0.0/index.html

[Y1731] International Telecommunication Unit (ITU), *ITU-T Y.1731 – OAM functions and mechanisms for Ethernet based networks*.

# 8. ANNEX

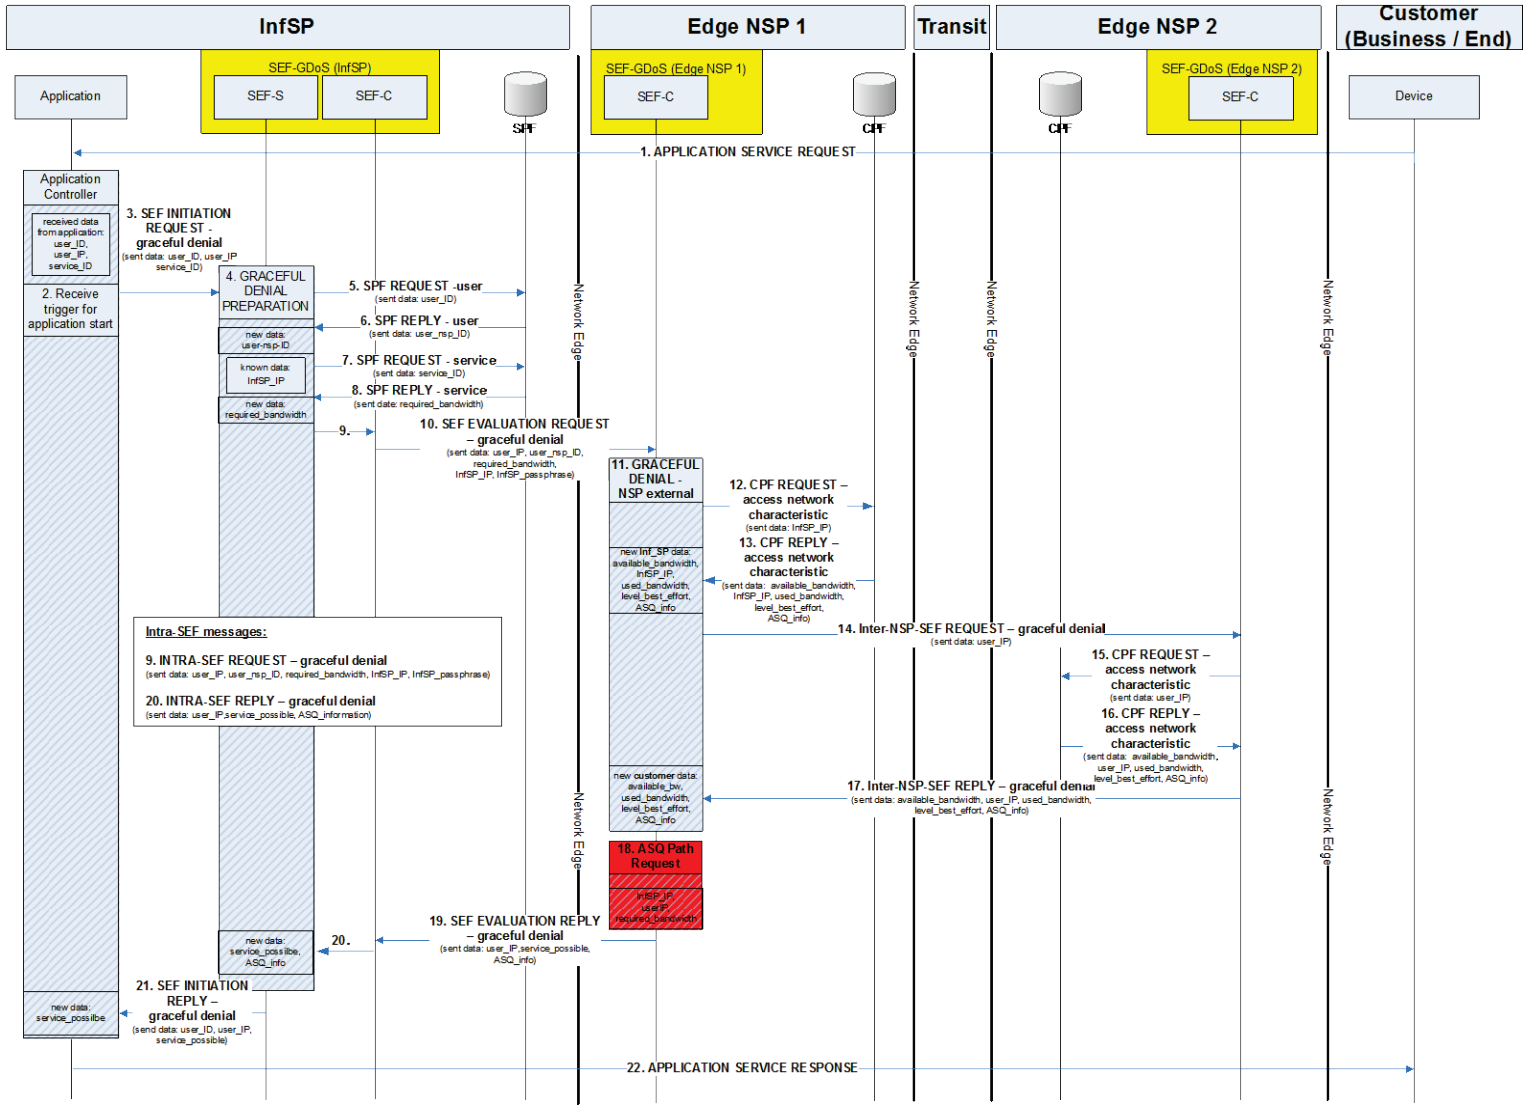


FIGURE 69: SEF-GDOS FLOW DIAGRAM