AGAVE

A liGhtweight Approach for
Viable End-to-end IP-based QoS Services

IST-027609

D1.1: Parallel Internets Framework

Document Identifier: AGAVE/WP1/FTRD/D1.1/public	
Deliverable Type: Report **Deliverable Nature:** Public	**Contractual Date:** 31 August 2006 **Actual Date:** 08 September 2006

Editor:	Mohamed Boucadair and Bruno Decraene, FTR&D
Authors:	*TID:* M. L. García Osma, A. J. Elizondo, J. Rodríguez Sánchez *FTR&D:* M. Boucadair, B. Decraene, B. Lemoine *Algo:* E. Mykoniati, P. Georgatsos *UCL.uk:* D. Griffin, J. Spencer; J. Griem *UniS:* N. Wang, M. Howarth, G. Pavlou, S. Georgoulas *UCL.be:* B. Quoitin
Abstract:	This document specifies the framework and architecture for designing, building, configuring and operating Network Planes within individual IP Network Provider domains and describes how Network Planes may be extended across multiple provider domains to form Parallel Internets suitable for meeting service-specific requirements. The business roles and relationships pertinent to the deployment of end-to-end IP-based QoS-enabled services are addressed and representative services are studied as business cases. Finally, this document includes the functional architecture defining the functions required for building the AGAVE solution.
Keywords:	Network Planes, Parallel Internets, QoS, Business Model, VoIP, VPN, Service Requirements, Functional Architecture

Telefónica Investigación y Desarrollo	TID	Co-ordinator	Spain
France Telecom Research and Development	FTR&D	Partner	France
Algonet SA	Algo	Partner	Greece
University College London	UCL.uk	Partner	UK
The University of Surrey	UniS	Partner	UK
Université catholique de Louvain	UCL.be	Partner	Belgium

Project funded by the European Community under the
"Information Society Technology" Programme (2002-2006)

Executive Summary

The ultimate goal of the AGAVE (*A liGhtweight Approach for Viable End-to-end IP-based QoS Services*) project is to solve technical problems related to end-to-end provisioning of QoS-aware services over IP networks, and to propose lightweight solutions compared to some existing proposals. The project approaches this problem by studying, developing, and validating an inter-domain architecture based on the novel concept of *Network Planes*, which will allow multiple IP Network Providers to build and provide *Parallel Internets* tailored to end-to-end service requirements.

This document addresses the problem of service provisioning and delivery across the Internet from the standpoints of both Service Provider and IP Network Provider. An interface is defined to enable their smooth interaction for the provisioning of the IP connectivity required by the IP-based services. As a means to accommodating services with diverse IP connectivity provisioning requirements running over a common IP infrastructure, the document outlines the Parallel Internets framework and architecture.

Specifically, a *Business Model* is drawn to capture the business roles and relationships pertinent to the end-to-end deployment of IP-based QoS-enabled services. A clear distinction between the roles of Service Provider and IP Network Provider is adopted, creating a business opportunity for specialised service providers to control without owning the IP infrastructure, and a new stream of revenue for the IP network providers who actually own it. Acknowledging the requirement for global coverage at both the IP and the service layers, providers are associated by horizontal interconnection relationships.

IP connectivity, Voice over IP and Virtual Private Network services are studied as *Business Cases* to set the requirements for the Parallel Internets framework and to drive the AGAVE work. Current business practices, latest developments, desirable enhancements and associated issues are investigated for each of the selected services. The *Requirements* of customers and service providers specific to each selected business case are envisaged.

A *High-level Specification* of the *Interactions* among service actors is provided focusing on the IP connectivity provisioning interactions, necessary to fulfil the derived requirements. IP Network Providers interact with each other to expand the scope of their offerings. Service Providers interact with IP Network Providers to control the provisioning of the IP connectivity portion assigned to their services. The interactions at the service layer are elaborated for each selected business case to demonstrate and validate the IP connectivity provisioning interactions.

The *Network Plane* concept is introduced and aspects of the Network Plane engineering process are investigated. Network Planes are brought into play for accommodating traffic with different end-to-end QoS and resilience requirements, coming from different service providers operating different service types. Network Planes are interconnected to build *Parallel Internets*. Alternative techniques for Network Plane and Parallel Internets realisation are presented.

A *Functional Architecture* is specified, covering the functionality required at both the service and the IP layer for the deployment of end-to-end QoS-enabled services following the paradigm of the Parallel Internets. Vertical and horizontal interfaces are defined between service actors. Service and IP layer functional blocks are identified and outlined.

Table of Contents

Detailed Table of Contents

List of Figures

1 INTRODUCTION

1.1 The AGAVE project

The overall goal of the AGAVE project is to solve technical problems related to end-to-end QoS-aware service provisioning over IP networks. The AGAVE project undertakes this effort by studying existing techniques and solutions as well as developing and validating a new inter-domain architecture based on the novel concept of Network Planes, which will allow multiple IP Network Providers (INPs) to build and provide Parallel Internets tailored to end-to-end service requirements. In light of the scope of this problem, the ambition of the project is to propose lightweight solutions that can be more easily deployed compared to some existing proposals. The project will specify an open connectivity provisioning interface to allow Service Providers to interact with underlying INP(s) infrastructure for the provision of their IP-based services.

1.2 Problem Statement

Despite significant efforts made by the research and development community to produce a set of IP QoS enabling mechanisms, protocols and architectures, QoS-centric techniques have not yet seen a large-scale deployment in operational networks. This is true for both intra- and inter-domain with some exceptions, as in the case of ISPs enforcing traffic shaping and policing rules at their access nodes and using marking techniques to distinguish flows at access segments while maintaining over-provisioning at the core of their networks.

In the meantime, new emerging services, such as online gaming and audio-visual streaming, as well as the migration of critical services such as PSTN telephony services, require IP networks to provide hard quality guarantees. These are guarantees in terms of availability of the service, for example to guarantee five nines availability of telephony services, to ensure emergency calls, as well as robustness guarantees to deal with "avalanche restart" and "flash crowds" phenomena, for example.

The AGAVE project believes that the current IP QoS approaches are incomplete. Experience has shown that existing QoS enabled frameworks and architectures only deal with specific aspects of the requirements of end-to-end inter-domain QoS services. Also there is no approach integrating requirements from both network- and service layers. This lack of consistency creates problems during the service (layers 5 to 7) design and deployment phase. The challenge of the QoS community and the AGAVE project is to bring existing and new ideas together with the inter-domain end-to-end QoS services requirements in a lightweight and easily deployable fashion, so QoS techniques may become standard tools for network operators' engineering practices and operations. This solution has to be based on the individual intra-domain mechanisms enforced within each domain as well as inter-domain mechanisms to enforce the end-to-end quality of the service.

Operators need to cooperate in order to offer consistent inter-domain traffic treatment to satisfy the end-to-end service requirements. Reducing the complexity of QoS enabled architecture will further promote and stimulate the cooperation between providers to offer inter-provider QoS-enabled services. One of the challenges of the AGAVE project is to develop lightweight techniques and architectures that stand alone as modules to facilitate the provisioning of inter-domain QoS services that are vendor independent and heterogeneous. This should simplify adaptation, deployment and operation of these new techniques, in terms of:

- Manageability of the proposed solutions, including Fault, Configuration, Accounting, Performance, and Security management;
- Service creation, offering, negotiation, delivery and assurance;
- Service evolution and maintenance;

The deployment of the proposed solution should take into account a critical factor, which is to keep deployment costs low, especially CAPEX and OPEX and to carefully elaborate and evaluate the migration operations and strategies.

1.3 Assumptions

As a starting point to the proposed work the AGAVE consortium assumes that:

- *IP is the universal transport network:* IP networks are the predominant transport networks for a large set of services like VoIP, CDN (Content Delivery Network), Games, etc., and the field for emerging applications such as peer-to-peer. The convergence with mobile networks (also known as Fixed-Mobile Convergence) is *not* in the scope of the AGAVE project;
- *Heterogeneity of services:* The services deployed upon IP networks are heterogeneous in terms of connectivity requirements, security support, sensitivity to delay and jitter, type of associated traffic (elastic or inelastic), traffic profile, etc. Hence, the transport requirements are different from a service to another one, leading to a requirement for differentiated treatment at the IP layer;
- *Complexity of Operations and Management*: As IP networks become the federative transport network, O&M related functions become (even more) critical. A single service may affect the operation of several protocols pushing the IP Network Providers to implement tools to ease the management of the services running over their networks and of the associated protocol operations;
- *Connectivity is more than reachability*: Offering IP connectivity can not be reduced to the single task of ensuring IP reachability towards a given destination, which can be enforced thanks to the activation of appropriate routing protocols. Additional constraints are to be taken into account when offering connectivity services such as providing reliable links and nodes with appropriate levels of redundancy, ability to support end-to-end security (authenticity, integrity, privacy) mechanisms (e.g., digestion, encryption), ability to cross clouds supporting a given protocol (like the support of RSVP), etc.;
- *QoS as a characteristic of the service*: Several techniques and algorithms proposed during the last decades are not introduced into operational networks due to multiple reasons, such as (1) the complexity of the proposed mechanisms, (2) the lack of clear views on the manageability of such mechanisms; (3) the reluctance of providers to abandon their practices related to over provisioning, etc. The QoS should be indeed studied not as a radical shift from the current practices but as an additional service feature relying on existing mechanisms, like appropriate tuning of the routing processes. The QoS is multi-dimensional and should not be treated only from one perspective. For instance, flow-aware routing can be used to differentiate treatment of IP packets but it is not sufficient if it is not associated with other resource management techniques;
- *IP QoS isn't about establishing LSPs everywhere*: Even if we believe that MPLS-TE (MPLS Traffic Engineering) is a QoS-enabler, (for instance, MPLS Fast Reroute allows guaranteeing sub-100ms recovery upon network failure), two misconceptions should be pointed out. First, LSPs (Label Switched Paths) do not ensure QoS nor give guarantees more than what is provided by other IP native means. Second, establishing LSP circuits between all nodes of the network can raise scalability concerns.

1.4 AGAVE technical approach

The AGAVE solution is built around the concept of *Parallel Internets* that enable *end-to-end* service differentiation across multiple administrative domains. Parallel Internets are coexisting parallel networks composed of interconnected, per-domain, *Network Planes*. Network Planes are established to transport traffic flows from services with common connectivity requirements. The traffic delivered within each Network Plane receives differentiated treatment both in terms of forwarding and routing, so that service differentiation across NPs is enabled in terms of edge-to-edge QoS, availability and also resilience. Within this overall scope the project will address a number of specific technical topics, including:

- Study representative business cases for the deployment of VoIP and VPN end-to-end services, derive requirements from the service provider and the IP network provider perspectives;

- Specify and validate an open IP connectivity provisioning interface allowing service providers to interact with IP network providers for the deployment of end-to-end services over Network Planes and Parallel Internets;
- Study alternative routing strategies such as differentiated routing and differentiated forwarding;
- Research new means to implement load balancing and sharing, especially inter-domain. This includes load balancing among several AS Border Routers, among several physical links or among several logical links (VLANs);
- Investigate means for Network Plane (NP) Engineering: network capability creation, NP definition and creation; inter-/intra-domain NP identifier management, inter-NP dependency, prioritisation of NPs, NP binding, extrapolation and extension options discovery;
- Research new techniques for failure management and routing convergence. The project may investigate how to extend emerging IP/MPLS Fast Reroute techniques, beyond domain boundaries (especially ASBR protection, PE ASBR protection, peering links protection).

1.5 Positioning this document

The AGAVE project is organised in four technical work packages:

- WP1, *Parallel Internets Framework*, is responsible for laying down the principles, concepts and reference architecture for realising the Network Planes and Parallel Internets and for enabling service provisioning from IP connectivity aspects adopting a clear separation of business roles;
- WP2, *Connectivity Service Provisioning*, undertakes the specification, design and implementation of appropriate interface components to enable end-to-end connectivity service provisioning and facilitate the deployment of IP-based services over the Parallel Internets;
- WP3, *Parallel Internets Engineering*, undertakes the specification, design and implementation of appropriate mechanisms, algorithms and protocols for realising the Network Planes and their interconnection over the Internet, to build Parallel Internets;
- WP4, *Validation and Experimentation*, undertakes the experimentation activities of the project for assessing the validity and cost-effectiveness of the proposed solutions.

This deliverable is produced by WP1, which is composed of two main activities.

AC1.1 Business Models and Service Requirements is responsible for developing the business models and specifying the service connectivity requirements. This activity identifies actors, business roles and interactions between Service Providers and IP Network Providers in the context of the deployment of IP-based services over the Parallel Internets.

AC1.2 Reference Architecture is responsible for specifying a reference architecture outlining the roles and the interfaces for engineering and providing connectivity services over the Parallel Internets based on the business models developed and the service requirements captured by AC1.1. The reference architecture and a set of appropriate accompanying use cases indicated by AC1.1 are used for defining the Parallel Internets framework.

This deliverable captures four principal components: the *AGAVE Business Model*, *Business cases definitions and requirements*, *High-level Specifications*, and *Functional Architecture.* Together, these form the formal output of WP1, concluding its activities. The models, definitions and high-level specifications documented in this report are a major input to drive the ongoing activities in WP2 and WP3. Detailed interface specifications, algorithm and protocol designs and finally implementations will be delivered by those WPs in due course.

1.6 Structure of this document

The rest of this document is structured as follows:

- Section 2, *Terminology and Definitions*, provides a list of definitions of terms used within this document, especially VoIP and VPN related terminology;
- Section 3, *AGAVE Business Model*, presents the business roles and relationships assumed by the project, distinguishing between the Service Provider and the IP Network Provider roles;
- Section 4, *Business Cases*, provides a description of the business cases that have been selected by the project to derive requirements from and to check the validity of its proposed solutions against. Three business cases have been identified and selected by the project (IP Connectivity, Voice over IP and Virtual Private Network services);
- Section 5, *Service Requirements*, lists a set of services requirements. These requirements are drawn from both customer and provider's standpoints for each one of the selected business cases. Three series of requirements have been elaborated: IP Connectivity services, (Inter-Provider) VoIP and (inter-domain/inter-AS) VPN ones. These requirements will be used by the project to derive IP connectivity provisioning requirements;
- Section 6, *AGAVE High-level Specifications*, provides an overview of Network Planes, their definition, creation to accommodate existing or future service requirements and techniques for their realisation. Additional issues, like interactions between service actors, are also provided in this section;
- Section 7, *Functional Architecture*, which describes the functional blocks required in order to build the AGAVE solution(s). The functional architecture spans the customer, service provider and IP network provider roles;
- Appendix A, *Issues in QoS*, discusses intrinsic QoS versus perceived QoS and the topic of Network Neutrality from a QoS and AGAVE viewpoint;
- Appendix B, *Terminology for resilience in IP networks*, provides definitions so as to reflect the notion of resilience, mainly within IP networks;
- Appendix C, *Overview of IP routing issues to support NP engineering*, discusses routing issues related to QoS and AGAVE viewpoint. Inter-domain, intra-domain and differentiated routing is discussed;
- Appendix D, *Inter-AS VPN*, is a state of the art of available standardisation documents on how to build inter-provider Virtual Private Networks.

2 TERMINOLOGY AND DEFINITIONS

This section provides a list of terms definitions as used within this document (see also section 10 for a list of abbreviations):

- Internet Telephony Administrative Domain (ITAD) denotes a telephony domain operated by a VoIP/telephony Service Provider. An ITAD can be located in a single AS, or span several ASes.
- Location Server (LS): this element is responsible for storing the location of customers which are connected to the service. LS can exchange its location information with adjacent LSes over a dedicated protocol.
- Virtual Private Network (VPN): A set of transmission and switching resources that will be used over a shared infrastructure to process the traffic that characterizes communication services between the sites or premises interconnected via this VPN.
- L3VPN: An L3VPN interconnects sets of hosts and routers based on layer 3 addresses (for example an IP destination address).
- L2VPN: An L2VPN interconnects sets of hosts and routers based on layer 2 identifier (for example an ATM VPI/VCI, FR DLCI, Ethernet VLAN, Ethernet MAC address).
- L1VPN: An L1VPN interconnects sets of hosts and routers based on layer 1 information (for example a port/interface, a wavelength).
- VPN Instance: From a management standpoint, a VPN instance is the collection of configuration information associated with a specific VPN, residing on a PE router.
- VPN Site: A VPN customer's location that is connected to the Service Provider network via a CE-PE link, which can access at least one VPN.
- VPN Service Provider: A Service Provider that offers VPN-related services.
- VPN Customer: It refers to a customer that bought VPN services from a VPN Service Provider.
- Customer Agent: It denotes the entity that is responsible for requesting VPN customer-specific information.
- Customer Edge (CE): Customer equipment like router or switch connected to the PE.
- Customer Premise Equipment (CPE): See CE.
- Provider Edge equipment (PE): Provider equipment connected to the CE.
- Provider equipment (P): Core router used to interconnect PE in the Service Provider backbone.

3 AGAVE BUSINESS MODEL

The business roles and agreements identified by AGAVE are illustrated in the following figure. It is important to stress that the identified actors described in the rest of this section represent roles which do not necessarily map one to one to distinct business administrations; a business administration may implement more than one role.

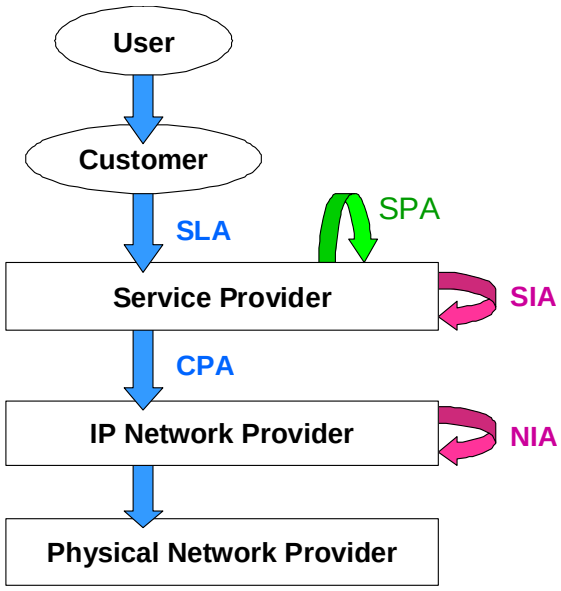

Figure 1 Roles and business agreements

3.1 Business Roles

3.1.1 Physical Network Provider

Physical Network Providers offer physical (up to the link layer) connectivity. Physical network providers are distinguished into two main categories according to their target market: *Facilities Providers* and *Access Providers*.

Facilities providers offer link layer connectivity to other providers in higher layers, *IP Network Providers* (INPs) or *Service Providers* (SPs) (see sections 3.1.2 and 3.1.3 respectively), for these to interconnect their own infrastructure and for interconnecting with their peers. Facilities providers may be further differentiated according to the technology they rely upon (e.g., optical fibre, satellite, antennas), deployment means (terrestrial, submarine or aerial) and their size in terms of geographical span and customer base.

Access providers provide the connection of the customer premises equipment to the INPs or SPs equipment. They own and administer appropriate infrastructure e.g. cables, concentrators. They may be differentiated according to the technology they employ e.g., POTS, FR, ISDN, xDSL, WLAN, Ethernet, as well as their deployment means and their size in terms of covered geographical area and customer base. Traditional state-owned PNOs in Europe and LECs in the US are typical examples of administratively distinct access providers.

3.1.2 IP Network Provider

IP Network Providers (INPs) offer IP connectivity facilities to *Service Providers* (see section 3.1.3). INPs own and administer an IP network infrastructure with the business objective to accommodate the IP connectivity requirements of their client SPs. INPs may provide Internet-wide connectivity or they may be active on a restricted area providing connectivity between specific end points.

INPs do not offer services to end *Customers* (see section 3.1.4). An Internet SP (ISP) associated with an INP provides IP connectivity services to end customers enhancing the IP connectivity offered by the INP with DNS, HTTP, e-mail service, etc. Note that INPs and SPs are roles, not necessarily distinct business administrations.

In fact the ISP term is commonly used to denote what in AGAVE's terminology is an ISP and an INP. As IP-based services proliferate, providers offering complete monolithic services give way to small specialised providers, which rely on outsourcing the IP connectivity provisioning to other providers to be viable. ISP/INP businesses on the other hand benefit from a new stream of revenue in opening their networks to other SPs to build their services over it.

Drawing the INP as an autonomous role interacting directly with SPs – from network layer SPs to higher layer Application SPs (ASPs) – introduces an interface between INP and SP, exposing the IP connectivity capabilities of the INP in a generic service-provisioning-aware but not service-specific way. This interface allows for multiple services operated by different SP administrations to run over a common IP network infrastructure transparently, with the INP optimising the network performance overall and under the constraints of each service running over it.

SPs interact with INPs following a customer-provider paradigm on the basis of respective agreements, called the *Connectivity Provisioning Agreements* (CPAs), regulating the IP connectivity requirements for the service control and data traffic, the control over and the feedback gained by the IP network operation offered by the INP to the SP.

For the purpose of expanding the scope of their IP connectivity, INPs interact with each other, on a one-to-one relationship basis. This interaction spans from the physical layer to the IP network layer (thanks to the widely deployed BGP protocol) for the purpose of exchanging "Internet full routing information" (subject to relevant routing policies) and is underlined by corresponding interconnection agreements, called the *INP Interconnection Agreements* (NIAs).

INPs may be differentiated according to the geographical span of their IP network infrastructure into small, medium and large INPs, with this distinction being relatively (to a given area size) rather than absolutely defined. For example, considering a continental area, small, medium, large INPs may be thought as regional (covering specific cities of a country), national (covering a specific country) and continental (covering specific countries of the continent) respectively. This distinction is essential from business perspectives, as INPs seek in augmenting the reachability of their offerings and is in line with current business practices.

3.1.3 Service Provider

The *Service Providers* (SPs) are responsible for the offering of IP-based services to the end *Customers* (see section 3.1.4). *IP Network Providers* (INPs) supply the IP connectivity over which the SPs set-up and offer IP-based services encompassing both connectivity and informational aspects e.g. telephony, content streaming services. As opposed to INPs, SPs may not necessarily own and administer an IP network infrastructure; they need to administer the necessary infrastructure required for the provisioning of the offered services e.g., VoIP servers and gateways, IP video-servers, content distribution servers. As such, for fulfilling the connectivity aspects of their services, SPs rely on the connectivity offered by INPs.

Note that the connectivity between the customer premises equipment and the service access point is in general the responsibility of the customer, which in turn establishes appropriate agreements with *Access Providers* (see section 3.1.1). In some cases however, an SP business administration acts also in the role of an access provider employing its own means for connecting customers to its infrastructure.

SPs interact with INPs on the basis of *Connectivity Provisioning Agreements* (CPAs). Beyond the forwarding and the QoS treatment of the service control and data traffic entering the INP's network from the SP's sites, the INP offers to the SP means to control the connectivity provisioning, such as receiving reports and alarms, invoking admission control at the IP network resources level, enforcing the required configuration on the INP's network on service activation, etc.

For expanding the scope and augmenting the portfolio of the offered services, SPs may interact with others. SPs may interact horizontally to expand the scope of one type of service, or vertically in a service provisioning hierarchy, where one SP enhances the service provided by the other SP to provide a new type of service. Such vertical business agreements between SPs are called *Service Provisioning Agreements* (SPAs) while horizontal business agreements between SPs are called *SP Interconnection Agreements* (SIAs). An example of hierarchically related SPs is in the provisioning of Web portals, where the portal SP enhances and bundles together services provided by other SPs, such as content distribution, online gaming, news SPs, etc. Note that SPAs refer only to vertical relationships between SPs for the purpose of providing a new higher-level service. Not every vertical relationship between two administrations that implement the SP role is regulated by an SPA; an SP business may be related to another SP business acting in the role of a plain customer (see section 3.1.4). A content service provider for example may be the customer of a telephony provider using the telephony service for its corporate network. In this case the relationship between the two administrations is a customer to SP relationship, irrespectively of the other business activities of the customer.

Different types of SPs may be distinguished according to the type of the offered services e.g., VoIP, ISPs, ASPs, Content Providers. SPs may be further distinguished according to their size in terms of customer base and/or geographical span, into small, medium and large (see above discussion on INPs).

3.1.4 Customer and User

A *Customer* is an entity, which has the legal ability to buy the services offered by a *Service Provider* (SP). Customers are the target recipients of the services created by the SPs. Services are offered on the basis of respective agreements, the so-called *Service Level Agreements* (SLAs), setting the terms and conditions on behalf of both SPs and customers for providing and requesting/accessing the services, respectively.

A *User* is an entity (human being or a process from a general perspective), which has been named by a customer and appropriately identified by an SP for actually accessing and using the services bought by the customer. The use of the services should be in line with the terms and conditions agreed in the SLA between the customer and the SP. In essence, users are the end-users of the services and they can only exist in association with a customer. A user may be associated with more than one customer using services according to the agreed SLAs of the respective customer. For instance, an employee may be acting as a user of the services that its company, customer, has subscribed to, as well as a user of its own subscription as a residential customer.

From the point of view of service provisioning, customers may be differentiated in terms of their size with respect to the number of geographical locations they may be present and/or the number of users they have, their type of business, the type of services they require, and the way and the habits in requesting and using the required services. Individuals, householders, small and medium enterprises, large corporations, universities or public organisations are typical examples of customers.

3.2 Focus of AGAVE

AGAVE focuses on the vertical business relationships between customers and SP and between SP and INP and on the horizontal business relationships between SPs and between INPs (see Figure 2). Business agreements between hierarchically related SPs will not be considered in AGAVE.

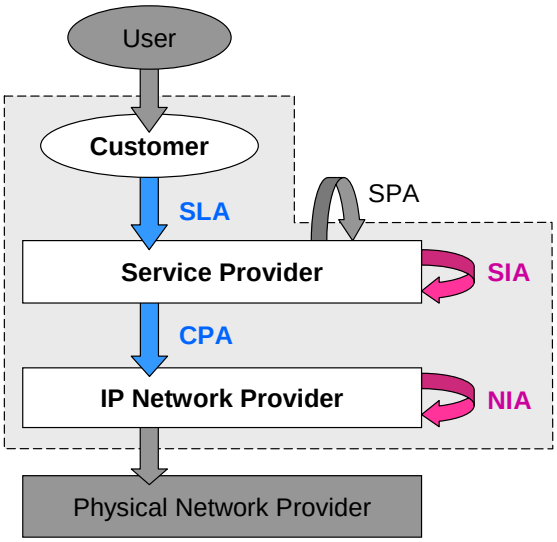

Figure 2 Scope of the AGAVE project

Interactions between the stakeholders will be primarily investigated from technical perspectives (specifications, protocols, etc.) with focus on the required functionality in SPs and INPs domains for handling and fulfilling the corresponding business agreements; accounting, pricing and legal aspects are outside the scope of AGAVE investigation.

4 BUSINESS CASES

In this section, we provide a description of the selected business cases. This description gives an overview of the service and identifies Layer 3 to layer 7 specific issues. A list of requirements of each service is elaborated in section 5.

Three business cases have been identified by the AGAVE project in order to check the validity of the proposed solutions. The project believes that these services represent the current and future mission critical applications.

4.1 Selection of AGAVE Business cases

Introducing the QoS in IP networks can be promoted thanks to the activation of some application drivers which may be (a) source(s) of significant revenue for Service and Network Operators. Nowadays, few applications are seen as QoS drivers, most of these applications are critical missions, especially because of the growing of bandwidth in core and access network. In addition, there are no heavy barriers for new market players due particularly to the openness of the technology and to the globalisation of knowledge. From this perspective, operators require demarcation arguments from each others and we believe that QoS will be one of killer arguments (of course the cost of the services is also to be taken into account!).

From this perspective, AGAVE project believes that its proposed solution(s) should be compliant with and suitable for these applications drivers so as to be adopted by operational entities. The project has selected the following business cases to check to validity of its proposed solutions:

- IP Connectivity services;
- Ensuring coherent QoS treatment across its (VPN) sites;
- Offerings robust and QoS-enabled Inter-provider VoIP services.

4.2 Role of Business cases

In order to drive the design and the specifications of the AGAVE solution(s), the project will rely on concrete business cases and cover the specific requirements of the corresponding services, at minimum as far as their IP connectivity aspects concerns. The selected service scenarios will act also as a reference for evaluating the pros and cons of the AGAVE solution(s) and other alternative solutions like the ones proposed by the IPSphere Forum (IPSF) [IPSF].

4.3 Description of the Business cases

4.3.1 IP Connectivity Services

One of the next challenges of IP networks is to introduce QoS-enabled IP Connectivity Services with a scope beyond the boundaries of a single domain and hence hopefully at the scale of Internet. Several techniques, architectures and protocols have been designed and promoted by standardisation bodies, like IETF, and implemented by major vendors, like Cisco and Juniper. Nevertheless the deployment effort has focused only on a single domain, managed by a single administrative entity, due to several reasons like the ones listed below:

- Lack of a clear business model for Internet wide QoS-enabled services;
- Lack of standardisation activities to harmonise and to drive the cooperation of several Service Providers (except the introduction of the notion of PDB);
- Lack of application drivers.

These aforementioned restrictions are not valid in the current stage especially with the migration of telephony services to IP infrastructures and the emergence of enterprise requirements like:

- Ensuring coherent QoS treatment across its (VPN) sites;
- Emulation of Pseudo-Wire lines spanning multiple domains [BRYA05];
- Connecting many islands of a single INP through the domain of another provider;
- Emergence of content providers requirements like VoD (Video on Demand) and other streaming service offerings like IPTV (IP Television);
- Etc.

QoS-enabled IP Connectivity Services can be divided into two categories:

- Limited expandability/restricted scope: For this type of services, the provider offers QoS reachability only to specific networks outside its domain. In this case, different QoS levels may apply to different networks. That is, a particular QoS level may only be experienced when reaching a specific destination network. This model is more aligned with enterprise needs.
- Unlimited expandability/open scope: The provider offers QoS reachability to any destination in the Internet, much like as today reachability is offered in the Internet at best-effort QoS levels. The offered QoS levels apply to all reachable destinations. This model is more suitable for the residential segment.

4.3.2 VoIP/ToIP Services

4.3.2.1 VoIP service definition

Telephony over IP is one of nowadays critical applications that succeeded to federate a large number of researchers and engineers in both standardisation bodies and industrial fora focusing on service providers' concerns. It is a mature technical area. As an outcome, several protocols and architectures have been proposed in order to deploy IP-based telephony service offering. A large part of these protocols are introduced progressively in operational platforms. In parallel, new service offerings for residential customers start to be deployed and promoted. Several challenges like the migration to IPv6 and the interconnection of providers are still open issues and should be solved by the community.

Within AGAVE project, a focus will be put on the inter-provider issue aiming to offer global and universal VoIP/ToIP communications. Within this document, ToIP denotes IP-based telephony systems which meet regulatory constraints like emergency, data retention and legal interception (LI). Excepting these regulatory issues, we believe that the placement of IP-based calls using VoIP or ToIP systems can be handled by using similar techniques and protocols. In the rest of this document, the term VoIP and ToIP are used interchangeably. Therefore, a set of functions may be supported (i.e. there is no constraint on the combination of the function that must be supported). A non comprehensive list of supported functions is the following:

- Location Service: The location service provides the mapping of the VoIP user identifier (commonly denoted as URI (User Resource Identifier) [BERN98]) to one or more addresses;
- Signalling: The VoIP SP undertakes and guarantees the delivery of all signalling communication from the point where the user call establishment request arrives inside the SP domain and beyond;
- Media Stream routing: The VoIP SP undertakes and guarantees the delivery of the call data flows;

4.3.2.2 VoIP service functions

The specific bundle of VoIP primitives offered by the VoIP SP is formed based on the business objectives of the SP and the needs of its target customer group. The primitives a VoIP service offering is decomposed into are the following:

- *Location Service*

The location service provides the mapping of the callee identifier to its current address, assuming the customers can be found in a group of addresses, some fixed and some dynamic. The location service is further decomposed into the directory service, presence service and translation service primitives, and can be offered as a whole or partially.

The directory service primitive provides retrieval of a list of telephony identifiers for the callee, e.g., the callee's home PSTN number, GSM number, Skype user name, etc. based on the callee identifier or any relative keywords, such as the callee's name, address, etc. The presence service provides in addition availability of the callee at each of the corresponding telephony identifiers. The translation service provides per telephony identifier the address of the next node to send the call establishment signalling message to, it is therefore signalling protocol aware, it will return for example the H.323 or SIP speaking gateway for the telephony identifier.

The presence service assumes some form of current location registration for the callee. Alternatively, a call forwarding plan uploaded by the customer for managing incoming calls could be used when the location service is combined with signalling forwarding, i.e., when call establishment is forwarded by the location VoIP SP, not by the caller itself.

The translation service may be used once for fixed addresses but it is required when IP addresses are dynamically assigned to VoIP terminals.

- *Signalling Mediation and Routing*

  Signalling mediation provides the translation of signalling messages between the two parties of the call, when the terminal equipment or soft-phone application of the two parties implement different signalling protocols. Signalling mediation is required, for instance, for PC-to-phone calls and may be even required for calls among IP based terminals to translate between different VoIP signalling protocols (e.g., H.323 and SIP). The VoIP SP has to deploy the necessary signalling gateways to be able to interconnect with domains based on other technologies.

  The VoIP SP undertakes and guarantees the delivery of all signalling communication from the point in SP where the user call establishment request arrives and beyond. This is an optional service for IP based terminals, for which the soft-phone applications may very well rely on the TCP/IP connectivity for reliable delivery of the signalling messages.

- *Media Flow Guarantees*

  The VoIP SP undertakes and guarantees the delivery of the call data flows. The delivered media stream guarantees will be, at least, some form of enhanced best-effort, being desired to have hard delay and loss guarantees.

  As long as the VoIP media stream flows through the IP network, either the customer or the VoIP SP rely to the underlying INPs for delivering the media stream. Delegating the media stream guarantees to the VoIP SP instead of going directly to the INP may be beneficial to the customer either because the VoIP SP offers better prices, because it offers better quality, or because it offers wider VoIP destinations coverage for a given quality and price.

  To implement the QoS guarantees the VoIP SP may rely solely on its INP to reach all destinations or it may delegate in a cascaded approach the media stream to other downstream VoIP Service Providers.

4.3.2.3 Terminology and procedure

The following figure summarises the functions required to illustrate a VoIP/ToIP Service Provider domain:

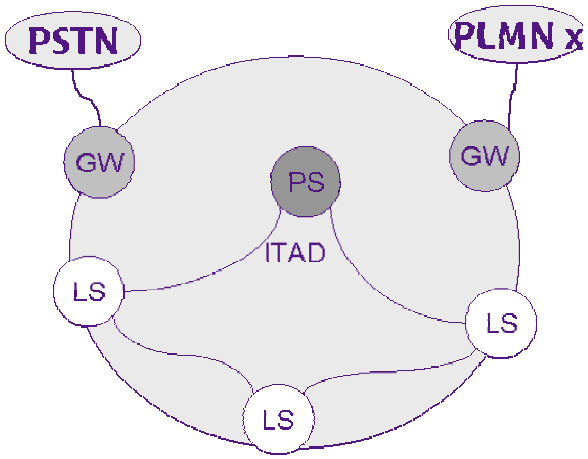

Figure 3 IP Telephony Administrative Domain

The scope of the VoIP Service Provider is denoted by the ITAD (*IP Telephony Administrative Domain*) [ROSE00]. The ITAD delimits the zone covered by the VoIP Service Provider and includes the Voice equipments it manages. In order to ease the representation of the VoIP Service Provider, the following elements/functions are introduced:

- PS (Proxy Server): This notion is similar to the one introduced in SIP architecture defined in [ROSE02]. Within the context of this document, PS denotes all required functions for the call routing and enhanced services supported by a VoIP Service platform.
- LS (Location Server): this element is responsible for storing the location of customers which are connected to the service. The LS is aligned with [ROSE00] but maintains also customers' registrations in addition to gateways capabilities. For instance, LS can be populated by a dedicated protocol such as TGREP (Telephony Gateway Registration Protocol, [BANG05]) or by subscription messages issued from user agents of authorised customers. The LS is invoked by PS only during the signalling phase.
- GW (Gateway): this element is used to interconnect the ITAD to external Voice platforms like the PSTN and other ITADs. Within the context of this document, the SBC (Session Border Controller) [ROSE02] is considered a GW. A GW intervenes during the media exchange phase and possibly during the signalling phase in case of presence of SBCs. This element is optional in case of interconnection between two ITADs deployed over IP but is mandatory, at least for trans-coding operations, between two heterogeneous telephony domains (e.g. PSTN and VoIP).

At the IP layer, these functional elements can be located in the same AS (BGP domain) (configuration 1) or in several domains (configuration 2) as illustrated in the figure below:

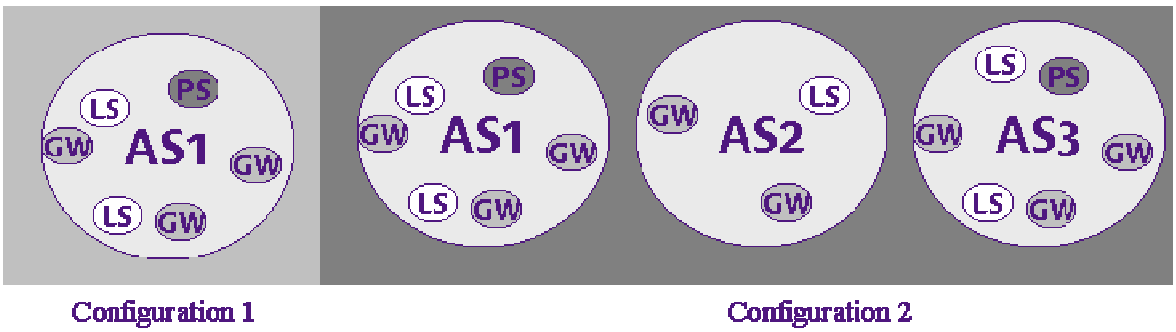

Figure 4 ITAD coverage

For instance, in configuration 2, the functional elements of an ITAD can be located in distinct BGP domains. These domains may or may not have a direct peering with each other. The cooperation

between the ITAD elements is managed at the service layer. In order to represent the Point Of Presence (POP) of a given VoIP Service Provider, the following scheme is adopted:

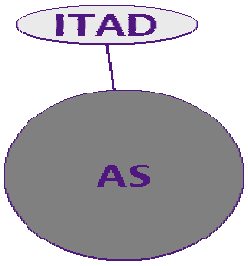

Figure 5 Relationship between an ITAD and an IP domain

The line drawn between the ITAD and the AS domain, represented in the figure above, expresses that IP connectivity of ITAD elements/nodes is ensured thanks to AS resources.

To benefit from the VoIP/ToIP service offered by an ITAD, customers and gateways register into the LS the IP address and port number they can be reached from (additional information can be enclosed in the registration request). The LS stores the location information if authentication procedure concludes successfully. In order to place a call, a customer sends its request to the PS which resolves the location of the callee by questioning its LSes. The call request is forwarded to the IP address/port number indicated by the LS. If the destination customer is not managed by the local ITAD, the PS routes the call request to an appropriate GW or to an appropriate PS. The following figure illustrates the roles of the ITAD and the IP infrastructure. The signalling messages are handled by the ITAD elements nevertheless; media streams can be exchanged directly thanks to the IP infrastructure without intervention of ITAD elements. Note that signalling messages flowing through the ITAD can be reduced to lookup requests, as the remaining ones could be directly exchanged between the involved edges.

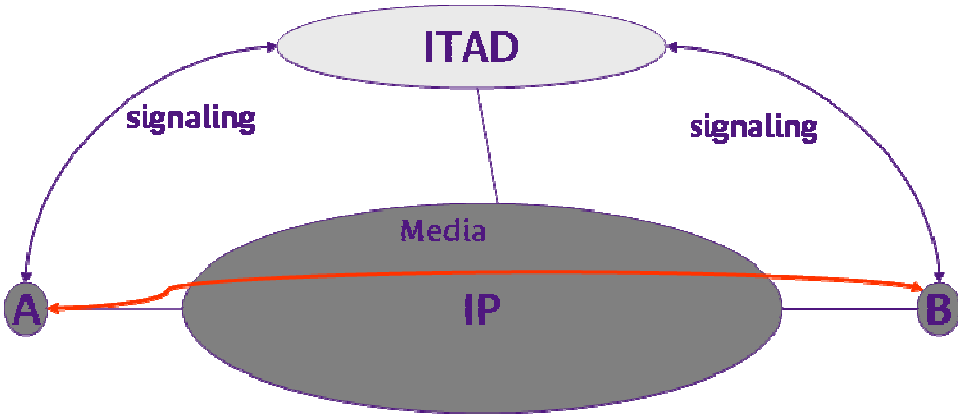

Figure 6 Role of the ITAD and the IP domain

4.3.2.4 Taxonomy of VoIP SPs

4.3.2.4.1 VoIP service capability options

As previously discussed, VoIP services consist of three broad sub-services: location service, signalling control and media stream routing. A VoIP SP may be involved at one of three levels, according to this functional decomposition. The lightest form of VoIP SP consists purely of a location service capability. In this case the SP maintains databases necessary for mapping number/id to IP address but is not involved in the establishment, routing or mediation of signalling or media flows. Signalling and media flows are established directly between users through the offered INP capabilities.

The next option for VoIP service provider capability is one that operates both location services and the control/mediation of signalling flows between users. The SP will mediate between different signalling

protocols (e.g., SIP<->H.323), relieving customers from the need to ensure their applications support multiple signalling protocols. This option allows SPs to enforce business relationships between themselves by controlling the routing of signalling messages through the sequence of SPs that have made prior agreements for termination and transit of VoIP signals. Note that in this option the SPs do not control the establishment and routing of media flows, this is done through direct interactions between customers and INPs.

The third option is where a VoIP SP includes all three levels of service capability, including the control of the establishment and routing of media flows. This is the only option where the SP has direct influence on the QoS associated with the VoIP call.

4.3.2.4.2 Centralised and distributed VoIP service provision

Voice and Video-based services emerge as critical services transported by IP networks mainly thanks to the event of the PSTN migration to full IP and the proliferation of new voice offerings like PC-to-PC, PC-to-Phone and Phone-to-Phone services. Nowadays, two types of VoIP business cases can be distinguished:

- (1) *The VoIP Service Providers which rely on the physical infrastructure to offer telephony-based services*. Historical PSTN operators migrating to IP infrastructure are part of this group. For this type of providers, the VoIP service engineering is closely associated with the underlying IP architecture. For instance, the QoS can be ensured by engineering dedicated VCs for transporting voice traffic at least between the subscriber home gateway and the first INP provider equipment or to assign dedicated IP addresses for telephony services (these addresses are not used beyond telephony services). This close interference between the service layer and the underlying IP infrastructure makes it difficult to offer the VoIP service beyond the boundaries of the INP invoked by the VoIP Service Provider. The cooperation of this category of VoIP Service Providers is driven by the physical infrastructure and the IP Connectivity offered by the underlying INPs. These VoIP Service Providers are subject to several regulatory constraints like ensuring emergency calls, legal interception and Number Portability. In addition, these SP are said to belong to the "*Centralised VoIP/ToIP Group*" since they (will) deploy centralised architectures following ETSI TISPAN and 3GPP IMS models. The cooperation of theses VoIP Service Providers would require the cooperation of underlying INPs.
- (2) *The VoIP Service Providers which offer voice-based services independently of the (IP) location of the subscribers*. This category of Service Providers does not rely on any agreement with the IP Network Providers. In addition, the Service Provider is not considered as an operator as it does not offer telephony services and is not obliged to meet the regulatory conditions as for the first case even though it provides exchange of Voice and Video between end users. The required investment to offer this service is drastically reduced. Note that the perceived Quality of Service depends on the capacity of the lines owned by the subscribers and their IP providers. Currently, there are two ways to offer this type of VoIP Services:
  - (2-1) The VoIP Service Providers which rely on a centralised service platform like Yahoo!, MSN, openWengoo and Damaka. At the service layer, these service offerings are considered part of the "*Centralised VoIP/ToIP Group*". For these services, the subscribers can benefit from the service offerings if reachability conditions are met. Some advanced services, like calls to PSTN realms can be controlled by the VoIP provider and to be offered only to a part of subscribers thanks to an enforcement of an access control based on the location (i.e., IP address) of the subscribers. These service offerings are achieved thanks to the collaboration of several servers in order to support big amount of load/traffic. The cooperation between these servers should be configured in order to offer consistent service offerings.
  - (2-2) P2P-based VoIP service offerings like Skype and Gizmo which do not rely "heavily" on central nodes, except for authentication before accessing to the service and "hard" configured super nodes. We refer to this category in the remaining part of this document as "*Distributed VoIP/ToIP Group*" (or Direct communications). The enforcement of this type of VoIP service is distributed mainly for lockup functions which are implemented by the nodes themselves. A

hierarchy can be brought in order to reduce the amount of operations handled by a given node and in order to ensure the availability of the service, even if this hierarchy is against the spirit of P2P architectures.

4.3.2.5 VoIP SP interactions

4.3.2.5.1 LS interactions

The interconnection between two ITADs is no more than allowing access to destination numbers not managed by a single VoIP/ToIP Service Provider. In other words, the local LS can resolve the locations where to send a call request even if this destination is not attached to the local ITAD. To do so, VoIP Service Providers need to exchange/share the location information, or at least to provide guidelines for the call routing logic of each SP so that calls to destinations not attached to the local ITAD can be placed successfully. This interconnection of ITADs and exchange of call routing information can follow two approaches:

- DNS-like approach: This mode is inspired from the DNS paradigm.
- Flooding approach: In this mode, two ITADs can peer with each other and exchange adequate information to guide the Call Routing Logic. This can be done statically or by activating a dedicated protocol which aims to propagate/discover and select an ITAD path for call establishment.

4.3.2.5.1.1 DNS-like approach

In this mode, the collaboration of ITADs in order to offer inter-provider VoIP calls does not require any exchange of information between ITADs. Each ITAD configures its DNS server with relevant information for the location of telephone number. The administrative registration and delegation of number prefixes is not detailed hereafter. An example of such architecture can be found in [FALT04]. In order to place a call, the callee telephony domain identifies the caller domain server. This is conditioned by the existence of appropriate entries in the DNS system. If no entry exists, the call can not be initiated. The following figure describes an example of the establishment of an inter ITAD call. Only signalling messages are represented.

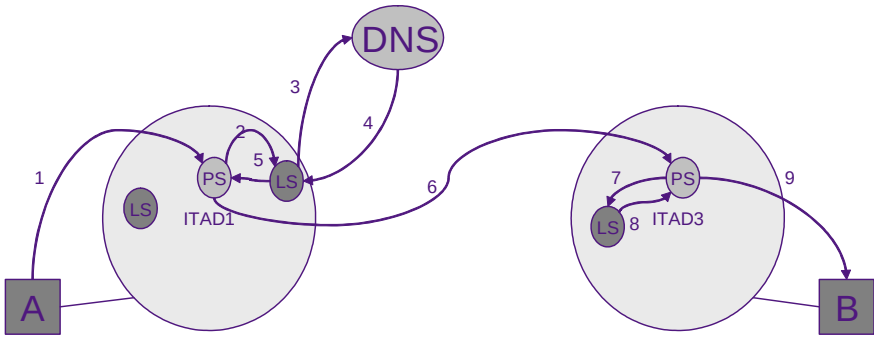

Figure 7 Example of call establishment

Hereafter some details explaining the messages drown in the figure above:

- (1) This message is sent by A in order to initiate a call towards B. This message is destined to the server PS responsible for call routing within ITAD1
- (2) Upon the receipt of a call request, PS invokes its LS in order to resolve the location of the caller B.
- (3) Since B is not registered within ITAD1, LS contacts the DNS in order to find the location of B.
- (4) The DNS provides LS with the appropriate information to reach the PS of ITAD2.
- (5) LS relays this information to its PS.

- (6) The PS of ITAD1 forwards the call requests to PS of ITAD2.
- (7) Upon receipt of the call request by PS, this latter invokes its LS to retrieve the location of B.
- (8) LS answers to PS with appropriate information to contact B.
- (9) Finally, PS of ITAD2 forwards the call request to the location provided by its LS.

Note that ITAD1 and ITAD2 can be attached to distinct ASes which are not directly connected.

4.3.2.5.1.2 Flooding approach

4.3.2.5.1.2.1 Principle

This mode is similar to the interconnection of ASes in the IP layer. Each ITAD advertises to its adjacent ITADs the prefixes it can reach. This information can be local to an ITAD or based on other information received from other peers. The following figure illustrates this behaviour:

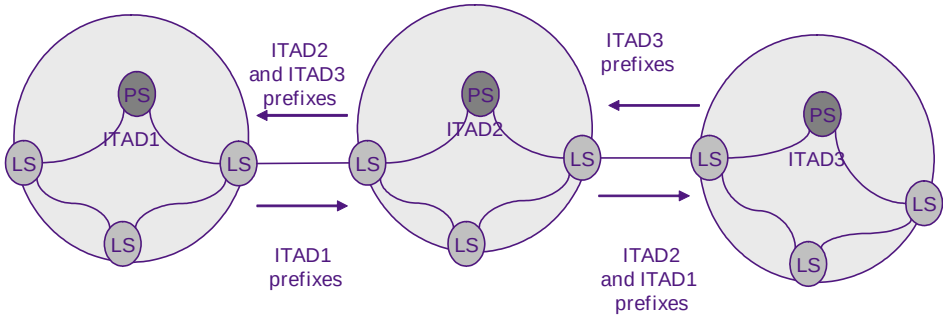

Figure 8 Flooding Approach

The Prefixes exchanged between adjacent ITADs can be static or dynamically updated by the inter-ITAD Call Routing Information exchange protocol.

4.3.2.5.1.2.2 Signalling interactions

4.3.2.5.1.2.2.1 Signalling Layer

When the flooding approach is adopted to exchange Call Routing Information between ITADs, two scenarios can be considered to set up calls towards destinations attached to a different ITAD as detailed below. Note that the procedure of exchanging, propagating and maintaining routes is similar in the two cases. The difference between these two scenarios resides in the manner the call is established between two parties each connected to distinct ITAD and not the way the telephony routing tables are built.

- Scenario 1: When a caller S wants to reach a remote telephone number, the associated user agent sends its call set-up message to its administrative server (PS). This server consults its associated LSes which do not find the destination number in the local subscribers' database, but thanks to the messages exchanged with its peers, the LSes of ITAD1 finds a route towards the callee domain. In this scenario, the LS of ITAD1 sends a request (not the call request but another lockup request) to the LS of ITAD2 in order to retrieve the IP address where to contact the PS of the callee domain. This procedure is repeated until reaching the ITAD8 which provides the IP address of the PS to be contacted. If the routing information maintained by all the LSes of the ITAD chain is correct, the PS of ITAD1 gets from its associated LS the IP address of the PS of ITAD8 where it should forward the call set-up request (This procedure is represented by the messages 2 to message 9 in the figure below) as a result of the communication between LSes described above. This communication between the LS is not the same procedure enforced by LSes to exchanges available telephony routes. Only external messages are illustrated in the figure). Finally, the call set-up message is routed to the PS of ITAD8 and then to D. D sends back an accept/reject message which is in its turn routed back to the PS of ITAD1 (using the same connection for receiving the

call request) and back to the caller. The figure below illustrates the signalling messages exchanged to set up a call between S and D.

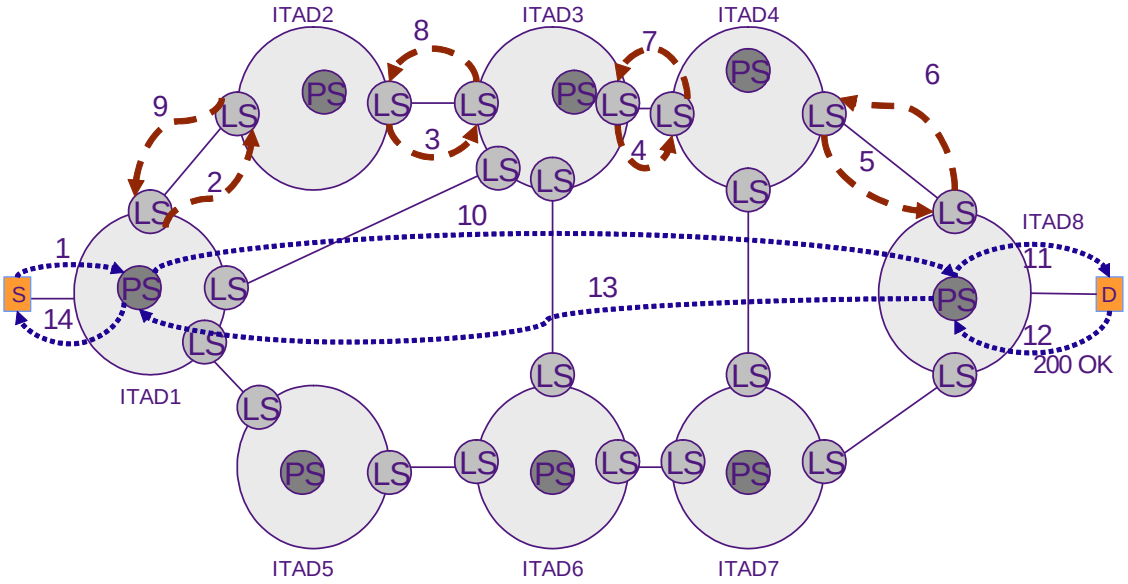

Figure 9 Inter ITAD Call Set up

- Scenario 2: When a caller S wants to reach a remote telephone number D, it sends a call set-up message to its administrative server, PS. This server consults its LSes which do not find the destination number in the local subscribers' database, but thanks to the messages exchanged with ITAD1 peers, LS finds an external route towards the domain managing D. LS of ITAD1 has also the next hop server to contact. This information is sent back to PS which forwards the call request to PS of ITAD21. This procedure is repeated until reaching the final ITAD to which D is attached. When the PS of ITAD8 receives the call request message, it invokes its LS in order to retrieves the location information of D. If D is connected to the telephony service offered by ITAD8, PS of ITAD8 forwards the call request to D. D sends back an accept/reject message which is in its turn routed back to the PS of ITAD1 (The path of the answer should be the same as the one of the request) and back to the caller. The figure below illustrates the signalling messages exchanged to set up a call between S and D.

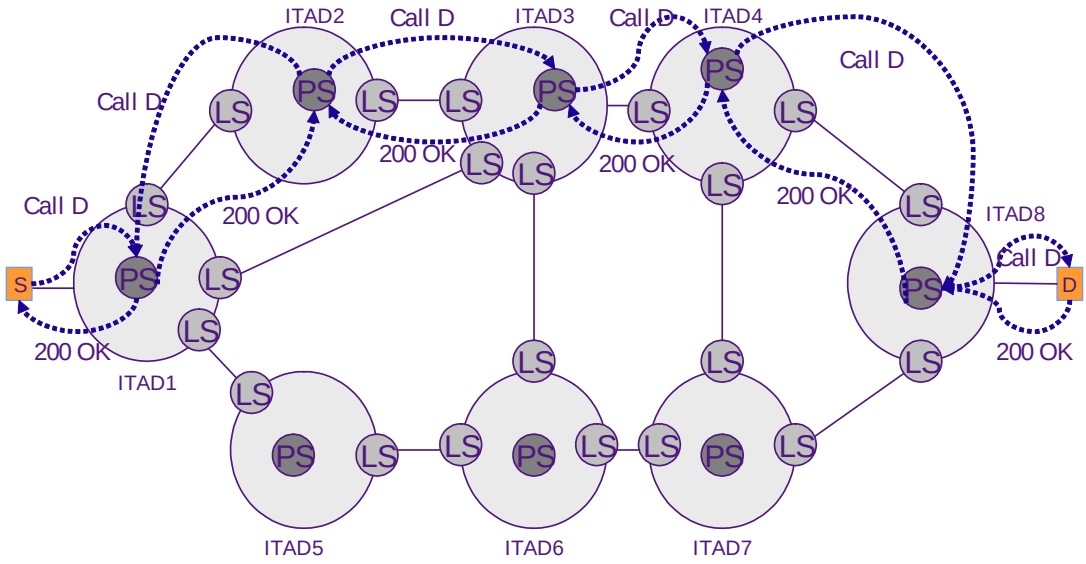

Figure 10 Inter ITAD Call Set up-bis

4.3.2.5.1.2.2.2 Media Layer interactions

After exchanging signalling messages between the callee and the caller, media sessions can be established and therefore media flows can be sent and received. The path followed by these media flows depends on the content of the signalling messages and the policies applied by the crossed ITADs. At least, two scenarios can be adopted. These scenarios are not a direct result of the two scenarios described above and apply for both them.

- Scenario 1: The media path is driven by the underlying INP route. Hence, it will follow the route determined by the corresponding originating INP to the terminating INP. The QoS treatment experienced depends on the selected path and crossed domains.
- Scenario 2: The media path is forced to cross the intermediate servers (SBC and GW). The followed path is forced by means of signalling messages which carry media session description information. In this scenario, the QoS treatment depends on the path selected in domains between the intermediate servers.

4.3.2.6 Distributed VoIP services

This section describes a VoIP service where all essential intelligence and service components are part of a client-side software application. Interaction with INP for network provisioning is minimised and the service is realised with traffic engineering in the overlay network. In the most extreme case there are no centralised hardware or software components. However, in most cases some centralised components may be present to improve performance and reduce complexity of the system. A distributed VoIP service is an overlay to an existing network that can span one to several INPs in scope or encompass the whole Internet and has to have location services as well as signalling components.

For distributed VoIP services the relationship between the service platform and the underlying IP network infrastructure is illustrated in the following figures.

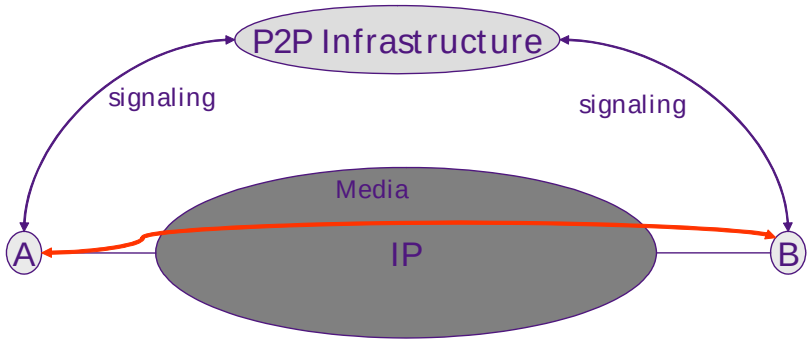

Figure 11 Example of P2P architecture

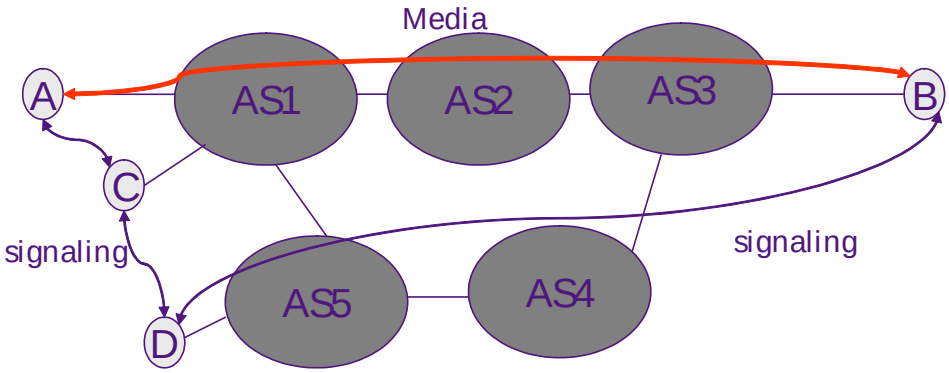

Figure 12 Example of P2P architecture-bis

Many proposals have been suggested and implemented for P2P architectures like Chord [STOI01] and Pastry [ROWS01]. Inter P2P platforms calls and P2P to centralised ITAD calls can have special requirements on the underlying IP infrastructure so as to deliver calls with the required quality and guarantees.

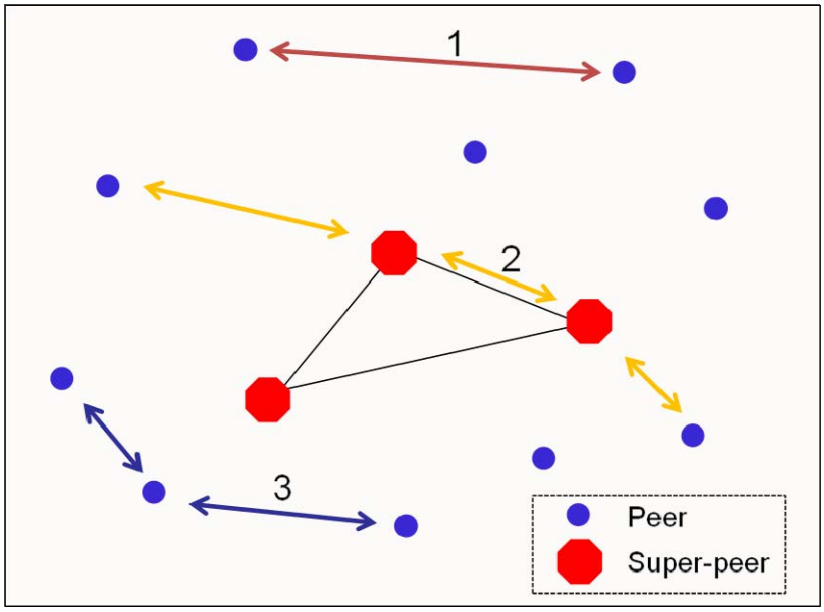

Figure 13 A Peer-2-Peer VoIP network with possible call routing scenarios

Figure 13 shows a typical P2P network where blue dots indicate ordinary peers and red dots indicate super-peers that may or may not be interconnected though dedicated connections. In case of a multi-plane Internet, interconnections could also be formed using premium plane services.

Overlay routing falls into the following categories:

- No overlay routing – arrow 1 on Figure 13
  - Direct point to point connection between calling parties
  - Direct point to point connection at premium plane service (if available)
- Routing via
  - Ordinary peers – arrow 3 on Figure 13
  - Positioned super peers – arrow 2 on Figure 13
- Routing using premium plane service via
  - Ordinary peers
  - Super peers

Category 1 does not require much intelligence in the overlay network, since no routing or node location awareness is required. Categories 2 and 3 are necessary if the direct best effort Internet connection between the peers is insufficient for high quality VoIP calls between two parties or if direct connections are not available because of firewalling or other network policy issues. Both cases 2 and 3 require more sophisticated overlay network intelligence to allow the network to select suitable relay peers that avoid network bottlenecks and have sufficient bandwidth and availability so as not to infringe the quality of the call themselves. Super-peers can be highly connected private peers as well as dedicated peers operated by the VoIP service provider in strategic locations. Third party super-peer ownership could also be conceived (such as belonging to access SPs). In theory super-peers are not necessary for the operation of the network, but it is presumed that they greatly improve service

performance and network stability. There is a soft boundary between the definition of a super-peer and a well connected normal peer, since no additional functions are performed by super-peers.

The functionality required to perform overlay routing is analogous to that of layer-3 network control and management operations. A simplified architecture of the functions required for operating a VoIP infrastructure is shown in Figure 14:

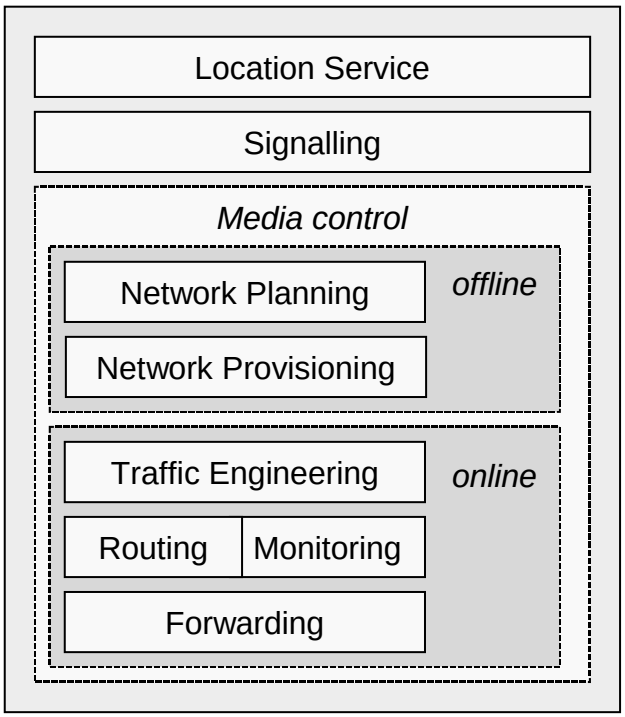

Figure 14 Overlay VoIP Service Components

The overlay network is created and maintained by the client-side software that runs the equivalent of the online components from Figure 14. This functionality is only required if routing via normal- or super-peers is a feature of the overlay VoIP service. Offline and service layer functionality is run by the VoIP service provider in a distributed or centralised fashion.

Routing is then responsible for finding suitable peers to relay via and Traffic Engineering is responsible ensuring that end-to-end quality is maintained. Since both functions are similar in nature, they are likely to be closely linked algorithms.

The Monitoring component is important for QoS management of VoIP calls, as it maintains feedback data for the relaying ability of peers and super-peers as well as measured network congestion, which aids Traffic Engineering in its choice of suitable relay peers.

If the overlay network is unable to create relays with sufficient quality two options remain to set-up a call. Both options escalate to the offline components, where the first is to request premium plane services to place either a direct or a relayed call with the premium plane service covering either part of or the entire route between the calling parties. This option, if available, could be requested on the fly or in bulk for individual VoIP customers through network provisioning. If this option is not available, or it is not cost effective, the second option is the creation of new super-peers managed by the VoIP service provider. This option is a network planning feature.

The terms “online” and “offline” are used here in the traditional network management sense. Some of the offline components could be requested on the fly, such as premium plane service requests as well as “upgrading” an existing peer to become a super peer. In general though, the components described as online are more “naturally” dynamic whereas offline services are invoked if the online services fail to achieve their goals.

Service layer functionality such as signalling and location services can also be implemented in distributed fashion, however, there is not much gain in doing so, since these services are low in bandwidth consumption and since databases are faster to access and update if they are centralised. Many existing distributed VoIP services therefore implement centralised location services.

4.3.2.7 Further considerations on Inter-Domain VoIP

4.3.2.7.1 INP Spiral

By INP spiral, we denote the act of crossing of a single AS several times in order to set-up a call. In order to illustrate this behaviour, let suppose that the ITAD chain to join D from S is {ITAD1, ITAD2, ITAD3, ITAD6}. Consequently the path followed by media streams is {AS1, AS2, AS1, AS6}. In this configuration AS1 is crossed two times (we refer to this as INP spiral).

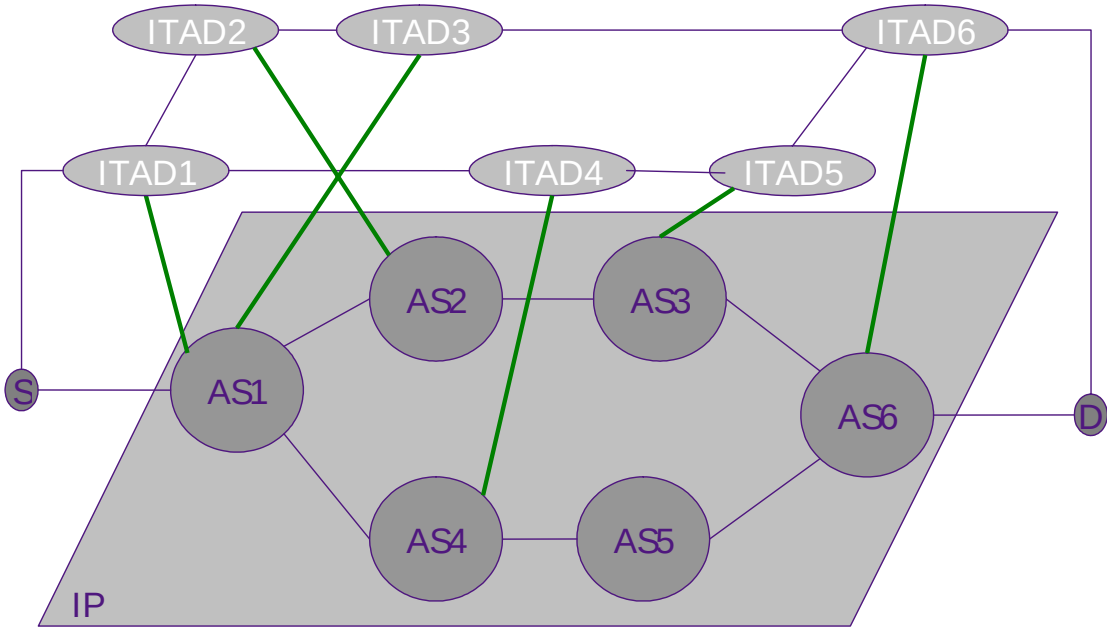

Figure 15 INP spiral

4.3.2.7.2 Synchronising the Control and Service Layer

In order to ensure the QoS for media flows established following the exchange of signalling control messages between adjacent VoIP Service Platforms, especially in case of scenario 1 detailed in 4.3.2.5.1.2.2.2, and to prevent to send media streams to domains that have not been crossed during signalling phase, the Call Routing and IP Routing should be synchronised.

In the example above, the media streams follow the path {AS1, AS2, AS1, AS6}. But, there is no inter-domain link between AS1 and AS6, so the crossed ASes are {AS1, AS2, AS1, AS4, AS5, AS6}. In addition, this path is not the optimal one to reach D.

4.3.2.7.3 No control of the INP path

Currently, the VoIP signalling protocols are not able to signal the underlying INP provider used to carry media streams neither to configure advanced policies like avoiding AS_PATH crossing a given INP domain. Indeed, two adjacent ITADs are not able to know during signalling phase, the INPs that will take in charge the transfer of media streams and whether they are using the same INP. This information can be used for path optimisation purposes like: selecting the next ITAD based on the underlying INP chain, avoiding spirals, etc.

4.3.2.8 QoS and inter-ITAD calls

4.3.2.8.1 Overview

The QoS is one of critical issues when dealing with voice over IP. Several initiatives have been launched within the IETF in this track especially those related to resource reservation associated with call set-up. As a result, a framework has been described in [CAMA02]. This proposal introduces a procedure for resource reservation during the call set-up phase. An enhanced proposal has been proposed by the EuQoS project [EUQOS], notably in the EQSIP proposal [VELT02] [VELT03] [SALS02].

4.3.2.8.2 QoS concerns

In order to evaluate the level of QoS experienced when placing a call, the VoIP Service Provider should evaluate the following information:

- End-to-end QoS parameters like one way delay, one way variation delay and loss rate
- Availability of the VoIP/ToIP service: mainly the availability of the service functional elements like gateways and proxies
- The nature of crossed links: whether those links are secured/backup, whether there are repair mechanisms that have been activated due to link failure, whether any crossed link is not, for instance, a satellite one, etc.
- Availability of bandwidth in order to place a call. Two approaches can be put forward:
    - Reservation per call: This option consists of negotiating the required QoS during the session establishment. The success of such a session is a necessary condition for the reservation of appropriate resources in both direction of the call. An example of this approach implementation is the QoS preconditions in [CAMA02]. This option has several drawbacks, being the main one that the connection set up time can be rather long if a call needs to cross several domains.
    - Service control access: Within this approach, the VoIP Service Provider does not reserve resources per call, but only verifies if the VoIP service platform can route the call based on appropriate information like the number of active sessions, the amount of supported simultaneous session, etc. In order to implement this option, adjacent VoIP Service Providers may set pipes (for instances, LSP or other forms of tunnels) between their server proxies or pipes involving a chain of ITAD domains. It is up to the domain initiating the call to inject the media traffic in the appropriate pipe based on the current state of the pipe.

Finally, it is recommended that inter-ITAD calls should be aligned with ITU recommendations related to international calls for both signalling and data transfer phases. Indeed, ITU has edited several recommendations related to telephony and associated performance targets. The following figure, extracted from E.820, illustrates the phases of a call and the related measures in order to evaluate the quality of the transfer.

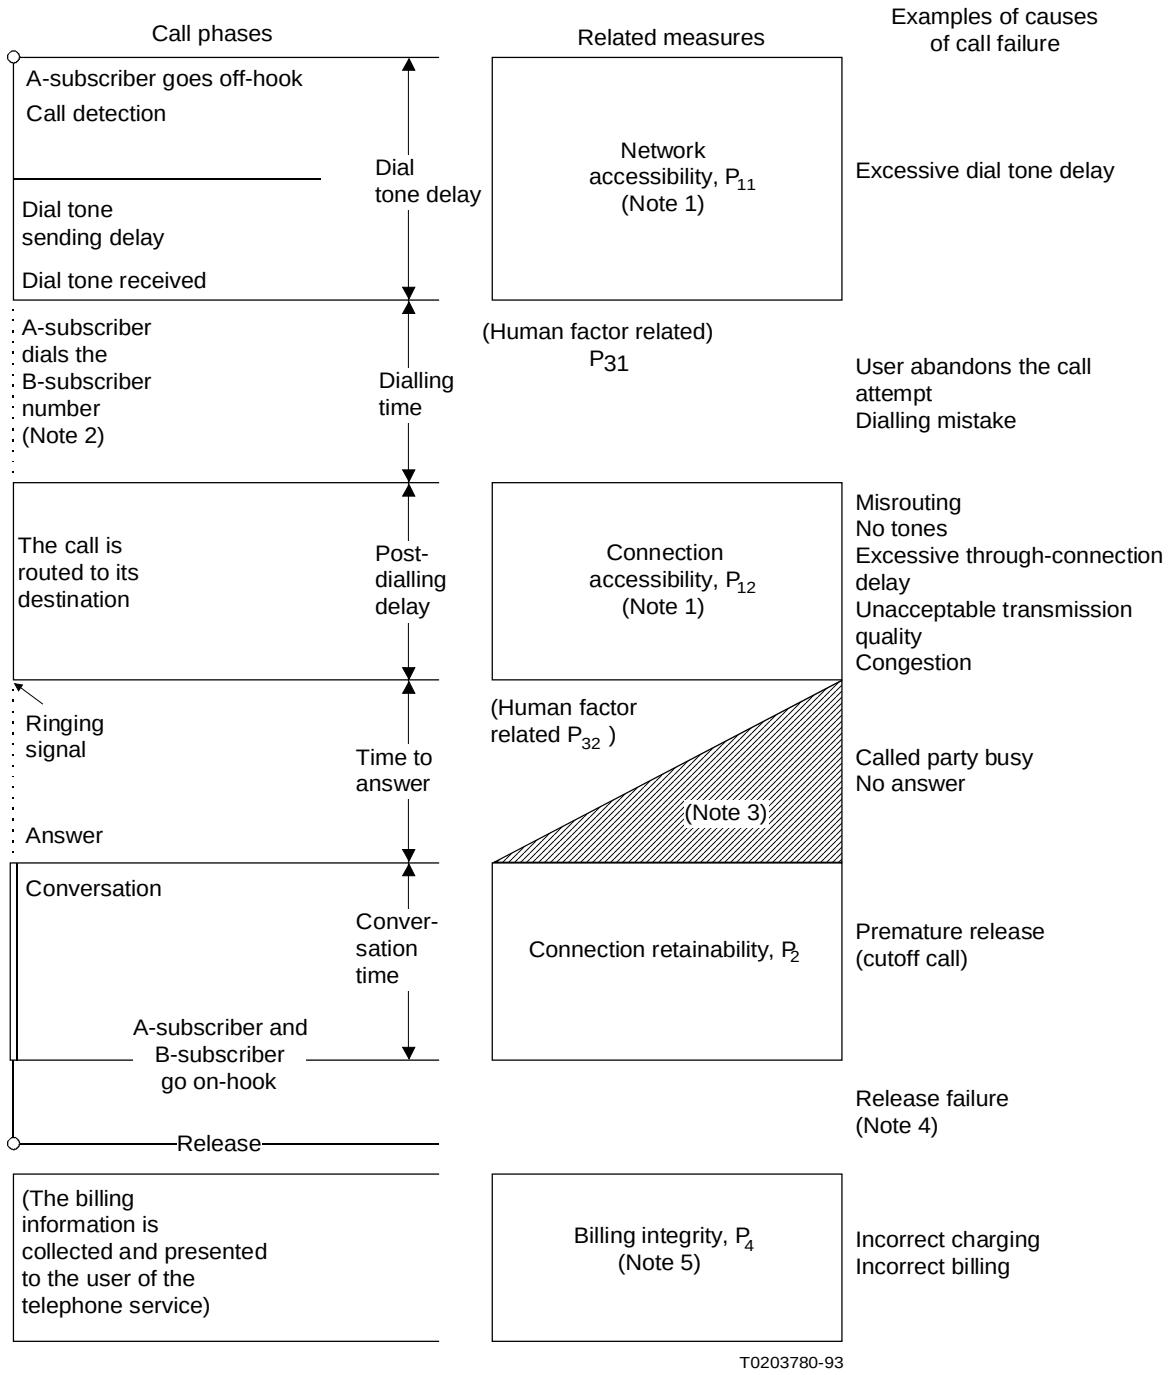

Note 1 – Network accessibility and connection accessibility combine into service accessibility.

Note 2 – The routing of the call may start before all digits have been received.

Note 3 – The shaded area shows that a premature release can occur during the time to answer.

Note 4 – The release of a call is not a separate phase in this model. A release failure may result in network inaccessibility for a new call.

Note 5 – The billing integrity has been shown for completeness, but is not a part of serveability performance.

FIGURE 1/E.820

Model of the serveability performance on a basic call in the telephone network

Figure 16 ITU Model of the serveability performance on a basic call

4.3.3 IP/MPLS-based VPNs

AGAVE focuses on L3 Provider Provisioned VPNs [CALL05] implemented using the BGP/MPLS approach [ROSE06], as this is the predominant model. The VPN state of the art with focus on interconnection considerations for the BGP/MPLS approach is presented in appendix 14.

A VPN customer uses the capabilities offered by the VPN SP to construct VPN(s) to interconnect its customer sites (intranets) and to further interconnect with sites external to the customer (extranets), e.g., under the administration of a partner organisation. The VPN customer interacts with a single VPN SP to construct its VPNs, irrespectively of the span of the VPN [NAGA04, CARU05]. From the customer point of view the alternative would be more complex to handle and would dilute SP responsibility. From the SP point of view this would reduce the added value and hence the profit margin. Hence, the VPN SP has to peer with other VPN SPs to reach remote customer sites attached to IP domains where the reference VPN SP has no point of presence. The association between two VPN SPs may be a peer partnership or it may follow the customer-provider. Relying on a transit VPN service provider under the cascaded model instead of establishing a direct peering relationship enhances scalability as the number of interconnections to maintain is reduced. Further, it may be the case that target QoS or resilience guarantees cannot be provided by the underlying INPs for the direct peering, while they can be achieved through the transit VPN service provider.

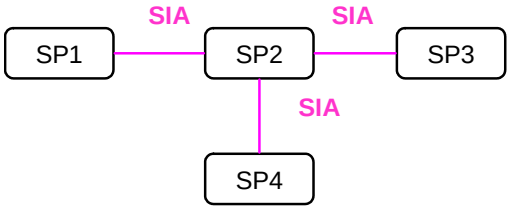

Figure 17 Cascaded relationship model

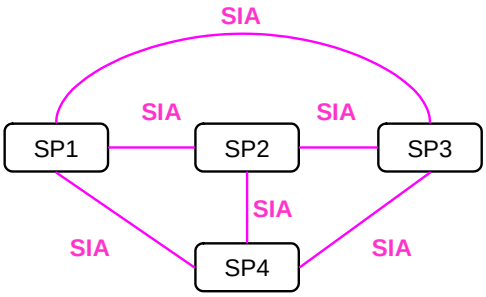

Figure 18 Bi-lateral relationship model

Finally, the VPN SP uses the connectivity provisioning capabilities offered by one or more INPs. An explicit requirement of VPN deployment is to limit VPN information to the edges of the network, keeping the VPN operation transparent to the core network equipment. The VPN SP uses the INP's core network and either employs dedicated equipment at the edges of the INP's network or it uses existing INP's edge routers either to attach customer sites or as VPN-aware border routers to interconnect a VPN AS with another VPN AS. Note that it is not expected to allow more than one SP to use a particular edge router of the INP.

For each presented scenario there is also its equivalent where SP1 and SP2 are not distinct but the same SP.

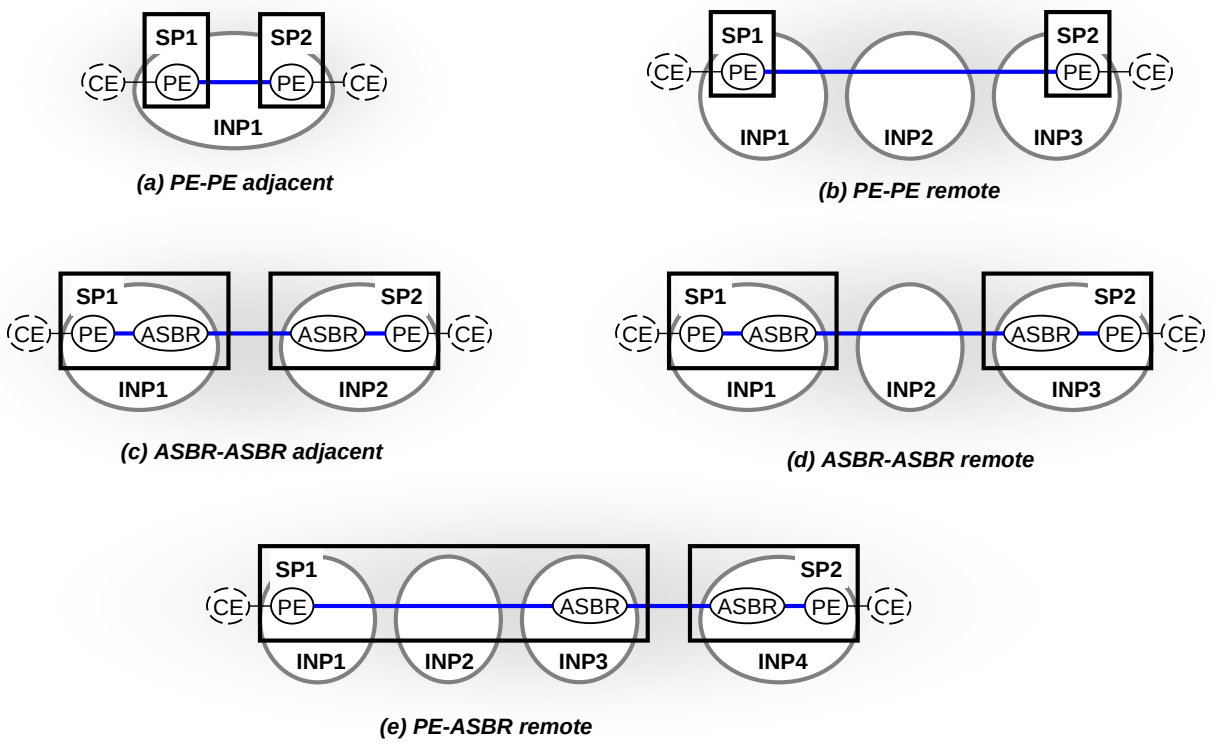

Figure 19 VPN scenarios

PE-PE adjacent scenario applies to topologies where VPN sites are attached to the same INP. PE-PE remote scenario applies to topologies where VPN sites are attached to non adjacent INPs and the corresponding VPN SPs rely solely on PEs for interconnection (see building inter-domain VPNs option "c" in section 14.2.3).

ASBR-ASBR adjacent scenario applies to topologies where VPN sites are attached to adjacent INPs and the corresponding VPN SPs deploy physically connected ASBRs for interconnection. Note that for this scenario and provided that the ASBR-ASBR link is entirely managed by the VPN SPs and not shared with INPs for other type of traffic, the INPs provide only intra-domain connectivity provisioning to the VPN SPs.

In ASBR-ASBR remote scenario VPN sites are attached to non adjacent INPs and the VPN SPs deploy ASBRs interconnected over intermediate INPs. See sections 14.2.1, 14.2.2 and 14.2.4 respectively for ASBR-ASBR inter-domain VPNs based on options "a", "b" or "d".

Deploying ASBRs to set-up VPNs improves scalability as tunnels interconnecting PEs are significantly reduced. PE-ASBR remote scenario applies to a VPN SP deploying ASBRs only in some of the INPs it operates over, and connecting PEs to these ASBRs over different INPs.

5 SERVICE REQUIREMENTS

Increasing the deployment of QoS-enabled services across IP networks and especially across the Internet requires a large set of providers to cooperate. This cooperation raises a number of challenges for providers due to the complexity of the technical issues to be solved and to the lack of appropriate standardized contractual agreements. To this aim, AGAVE will design, specify and validate a lightweight inter-domain IP QoS architecture and appropriate solutions to promote the deployment of QoS-enabled services.

The main purpose of this section is to identify a list of requirements, related to the actors involved in the QoS-enabled services delivery chain for a limited set of services like IP connectivity, VoIP and VPN services built over an inter-domain QoS-aware IP architecture.

The organization of this section is as follows. For each selected business case (see section 3), the set of identified requirements are grouped in two categories: requirements from the customer standpoint and requirements as seen by providers (both INP and SP are considered).

The customer requirements sections present requirements from the perspectives of the customers of QoS-enabled services built upon IP networks. The requirements are drawn from current business practices and market needs as understood by the project partners. The customer requirements pose corresponding requirements to the providers (both INP and SP), which in turn need to be taken into account into the IP layer solutions that will be proposed by the AGAVE project.

5.1 IP Connectivity Services

5.1.1 Customer Requirements

5.1.1.1 IPCS-C1: Definition and procurement of QoS guarantees

Customers should be able to freely choose their QoS-enabled services according to their actual needs. Thus, customers should be ideally offered with a choice of IP connectivity services at different QoS levels. However, when the service is actually requested, its QoS levels should be clearly and unambiguously defined.

The QoS guarantees should refer to well-defined performance metrics reflecting the QoS from the customer's perspective. At the network layer, these metrics should reflect the packet transfer quality (e.g., throughput, one-way transit delay, inter-packet delay variation -jitter-, and packet loss) and the availability of the connectivity service. The end-to-end QoS guarantees should be clearly specified, commonly understood and mutually agreed by the customers and the service providers.

- The QoS could be quantitatively specified, either statistically or hard guaranteed, by means of guaranteeing certain bounds on related performance metrics. These bounds are absolute for hard guarantees, in the sense that they cannot be transgressed. Statistical QoS means that they are normally fulfilled but with exceptions, the occurrence of which is quantitatively bound.
- The QoS could be qualitatively specified or soft guaranteed, e.g., relatively to other QoS levels by means of appropriate qualifications such as golden, silver, bronze QoS levels.

The above cases are distinguished because they refer to different types of customers, in terms of their requirements in using QoS services; therefore corresponding to different business cases. Some customers may know in advance the type of QoS services they require, whereas some others may not.

From a service provider's perspective the above requirements yield the following requirements:

- The SLAs underlying the offering of QoS-enabled services should:
  - Capture the QoS characteristics of both inbound and outbound traffic (with respect to the premises of a customer),
  - Specify the QoS characteristics quantitatively and/or qualitatively.

5.1.1.2 IPCS-C2: QoS topological scope

Customers should be able to send/receive traffic to/from any source/destination in a given pre-defined range. This range can go from a private network scope to the entire Internet scope.

In the private network scope, the customer should be able to:

- Send traffic with end-to-end QoS guarantees to specific destinations i.e., only to destinations, which have been *a priori* agreed with the service provider.
- Receive traffic with end-to-end QoS guarantees from specific sources.

In the Internet scope, the customer should be able to:

- Send traffic with end-to-end hard or statistical QoS guarantees for *a priori* agreed destinations or for dynamically requested destinations subject to availability, and with qualitative guarantees to any destination.
- Receive traffic with end-to-end QoS guarantees from any possible source with similar constraints as for sending traffic.

From the perspective of the IP network provider the above requirements yield the following requirements:

- To be able to expand the geographical span of the offered QoS services beyond the INP's domain.

5.1.1.3 IPCS-C3: Dynamic service subscription

Subscriptions should not be taken for granted as long-lived service contracts. Subscriptions may well be short-lived e.g. for a weekend. In fact, given the multi-service, multi-provider nature of the telecommunications market and the dynamic nature of customer needs -not all customers know in advance their QoS service needs-, the ability to establish SLAs dynamically is a key aspect of service offering. Some customers may be more attracted by such dynamic service offerings compared to static, monolithic offerings, as their service needs continuously evolve.

From a service provider's perspective, this requirement yields the following requirements:

- Providers should provide means for enabling customers to modify and terminate existing service contracts.
- Providers should provide means for enabling customers to subscribe to QoS-enabled services on-demand and for short time periods, upon customer requests.

Automated means for enabling subscription e.g. through Web servers and for handling subscription requests e.g. service configuration/activation means, could facilitate the satisfactory fulfilment of the above requirements.

5.1.1.4 IPCS-C4: Service invocation

Service invocation means the ability of customers to invoke i.e. to request to actually use the QoS IP connectivity services. Services are invoked by the users, within the invocation profiles agreed between the customer and the service provider at subscription time.

This requirement entails the following:

- Customers should be able to invoke the IP connectivity services either explicitly or implicitly. Explicit invocation will probably yield the use of an explicit signalling protocol. Implicit invocation does not require the explicit use of a signalling protocol; users can initiate their flows at any time, as long as the corresponding streams adhere to agreed subscribed profile.
- Customers should be provided with appropriate means to invoke their QoS IP connectivity services. These means should be in accordance with the service specifications.

From service provider perspectives, this requirement yields the following requirements:

- Providers should be able to support both explicit and implicit service invocation. As for the former case is concerned, providers should be able to support appropriate signalling protocols.
- Providers should provide for automated means in authenticating and authorising a (implicit or explicit) request of a QoS IP connectivity service.

5.1.1.5 IPCS-C5: Self-monitoring means

Customers should be able to check that the quality of the services they have subscribed to is in accordance with what they have agreed with the service provider. This requires that they should be provided with appropriate self-monitoring means.

From a service provider's perspective, this requirement yields the following requirements:

- Providers should provide customers with appropriate means, enabling the customers to assess the effective service throughput and QoS, to inspect their active inbound and outbound flows, etc.
- Providers should cater for appropriate customer care means for receiving and analysing customer complaints with respect to the service compliance to the agreed terms.

5.1.1.6 IPCS-C6: Protection against QoS (D)DoS attacks

The reception of undesired traffic may disturb the reception of other traffic, especially if the former is being received with higher QoS guarantees.

Due to the fact that QoS is, from the customer point of view, an end-to-end issue, all the participants in a given QoS-enabled session should agree on the QoS level of their communication.

The provider should block the undesired QoS traffic before reaching the customer by applying policing and/or remarking rules for the customer's inbound traffic. The rules are statically and/or dynamically configured by the customer. In the latter case the provider should provide means to the customers for dynamically configuring the policing and/or remarking rules.

Security is out of scope of the AGAVE project.

5.1.1.7 IPCS-C7: Multicast support

Since almost all the multicast services are receiver oriented, the following issues arise from the perspectives of multicast receivers (group members):

- Receivers desire to receive only interested multicast traffic from the subscribed group, and hence the functionality of source filtering is needed to avoid reception of unwanted multicast traffic.
- Receivers should be able to specify their QoS requirements individually, i.e. different recipients could specify different QoS levels via receiver-oriented SLAs for multicast traffic.
- Fairness among group members with heterogeneous QoS subscriptions of the same customer. For example, if one BE receiver is attached on the multi-access network where there is an EF receiver for the same group session, the BE receiver will receive EF traffic anyway even if a per-NP-tree approach is adopted.

Multicast is out of scope of the AGAVE project.

5.1.2 Provider requirements

5.1.2.1 IPCS-P1: Extension of QoS topological scope

The AGAVE solution should ensure that an IP network provider can offer QoS-enabled connectivity services, spanning beyond its domain i.e. across multiple ASes, with QoS levels coherent with the ones it is able to offer for intra-domain traffic.

This requirement breaks down into the following two non-exclusive cases:

- Limited expandability: The provider is able to offer QoS reachability only to specific networks outside its domain. In this case, different QoS levels may apply to different networks. That is, a particular QoS level may only be guaranteed when reaching a specific destination network.
- Unlimited expandability: The provider is able to offer QoS reachability to any destination in the Internet, much like as today reachability is offered in the Internet at best-effort QoS levels. The offered QoS levels potentially apply to all destinations.

5.1.2.2 IPCS-P2: Discovery of IP network providers and their capabilities

An IP network provider should be able to discover other IP network providers and their capabilities in order to easily and quickly determine the appropriate partners (from a business perspective) for expanding the scope of QoS connectivity. These capabilities may include the QoS and resilience guarantees provided per IP destination, the forwarding capacity, a list of the IP destinations directly attached to the provider, etc.

5.1.2.3 IPCS-P3: NIA flexibility

Once potential interesting QoS offerings have been identified, the process of reaching a common agreement should be rapid and easy. This means that the process for establishing interconnection agreements should follow accepted business practices, well-defined procedures involving finite steps based on commonly understood notions. Relevant automated means are desired for speeding-up the process.

IP network providers should be given means for requesting, negotiating, establishing, modifying and deleting these agreements. Caution should be taken to ensure that the modification of a given agreement does not disturb, but is in accordance with the requirements of other agreements related to the agreement under modification.

In case of agreement removal, means must be provided to ensure the coherence and stability of the system, notably the good handling and management of upstream agreements that were using resources provided by the downstream agreement to be removed (in a cascading approach). Possible solutions are for instance: forbid the deletion, notification to peers so that they modify their agreements before removal is completed, etc.

5.1.2.4 IPCS-P4: NIA and CPA assurance and monitoring

NIA and CPA assurance denotes the ability to check that what is being provided conforms to what has been agreed contractually. AGAVE solution(s) must enable inspecting the conformance of the effective interconnection and connectivity provisioning against the contractual expectations. Means must be provided to the INP for deriving assurance data from its own network concerning the handling of incoming traffic, and from the interconnected INPs for their handling of its outgoing traffic. Similarly, means must be provided to the service provider for deriving assurance data for the handling of its traffic by the underlying INP. The network configuration and policies derived by the AGAVE system must be in accordance with the QoS guarantees agreed in the contract.

5.1.2.5 IPCS-P5: Scalability

Within this document, scalability denotes the ability of the system to function effectively and keep its performance in the desired levels, as the size of the parameters influencing its behaviour increases. In other words, the proposed AGAVE solution should keep its performance unaffected whatever the size of domain span, which could be expressed in terms of number of participating domains (and routers), and whatever the number of agreements to be dynamically negotiated and invoked. The volume of the QoS-related information propagated across domains should be kept low, and should not affect the overall system performance, stability and (access) availability of the IP networks themselves. AGAVE solutions should cater for evaluating their scalability. This entails the assessment of the complexity of the decision-making components.

For instance, typical size parameters to take into account include:

- Per AS: average number of peers and average number of Network Planes.
- Globally: number of participant ASes, number of NIAs, number of Network Planes, and number of SLAs and CPAs to accommodate.

5.1.2.6 IPCS-P6: Resilience differentiation means

By resilience, we mean the ability for the system to recover from a failure by repairing itself automatically without restarting the service. Within AGAVE, this means among others that, in case of failure (e.g., link rupture or router breakdown), the system must be able to find/propose another path of equivalent QoS for the affected destinations. This operation must ensure that all active flows are automatically redirected correctly (i.e., no routing loops) with minimum disruption.

The level of resilience is evaluated with respect to the experienced disruption. To minimise disruption the IP network provider needs to proactively configure alternative paths provisioned with spare capacity to be utilised in case of failure. A different level of resilience may be required by different services and service providers, or even by different IP connectivity service customers. Hence, the IP network provider needs to have means to differentiate the level of resilience within and beyond its network.

5.1.2.7 IPCS-P7: Manageability

There are two main domains covered, which must be tackled by AGAVE, in this area:

- Configuration:
  - The base configuration, which is intrinsic to the solution, must be manageable and automation must be provided.
  - The configuration induced by the enforcement of a new agreement must not be too heavy, nor make the system unstable (not even shortly).
  - The impact of a system modification (for instance, a modification of an intra-domain QoS class) must be limited, and must not leave the system unstable (not even shortly).
- Monitoring:
  - The system must offer high-level monitoring information suitable for facilitating offline performance analysis, online processes, system operation and troubleshooting.

5.1.2.8 IPCS-P8: Backward compatibility

In order to achieve the goals pursued by AGAVE, proposed solutions are likely to introduce modifications on the existing infrastructures. The AGAVE approach should provide backward compatibility, not only to allow a smooth migration, but also to prevent existing infrastructures from being unusable and unstable.

Among other criteria, the following are considered as important to assess the fulfilment of this requirement:

- The impact on the intra-/inter-domain routing processes must be as limited as possible.
- When in operation, the AGAVE system must not introduce instability neither on the network itself, nor on the already deployed and running services.

5.1.2.9 IPCS-P9: Deployment easiness

The deployment easiness is related to how must time and effort it would require to have all the building blocks ready for operation, that is to say, to begin actual inter-domain communication with QoS activated. The easiness of deployment depends on a number of parameters, such as: number of

new protocols required, degree of adherence of the proposed solutions to the market and capabilities of commercially available routers, magnitude of required modifications to existing protocols, impact on intra-domain routing, impact on inter-domain routing and required conformance of other providers with the proposed solutions. The AGAVE solution(s) should clearly identify and describe such aspects.

5.1.2.10 IPCS-P10: Multicast aspects

It would be interesting to evaluate the impact of supporting multicast IP connectivity services on the features and performance of the approach along the following lines:

- Does the multicast support imply major changes or add-ons to the unicast model?
- How to manage the replicated multicast traffic within the network?
- How to avoid imposing significant impacts on the underlying IGMP, PIM-SM, MP-BGP protocols, as well as core router architecture for including QoS-aware multicast services?
- How to handle the scalability issues concerning QoS deployment?
    - Low overhead for group/QoS state maintenance within core networks.
    - No traffic conditioning at core routers.

Multicast is out of scope of the AGAVE project.

5.2 VoIP/ToIP Services

The purpose of this section is to identify the VoIP-specific (or generally related to conversational services) requirements in order to offer telephony services with world-wide coverage. For more information about the VoIP/ToIP use case, please refer to section 4.3.2. Note that a set of requirements have been identified and described in [MULE05], but some of these requirements are protocol specific and not generic enough so as to cover all VoIP provider interconnection issues. The requirements described below do not assume a specific signalling protocol between VoIP domains neither make assumptions on the way end-to-end (telephony) sessions are established (concerning signalling, security, media exchange, etc.).

In this deliverable, we assume that no single VoIP service provider can ensure global reachability to offer universal (telephony/conversational) services, nor has geographical presence to offer legacy telephony services globally. From this standpoint, cooperation between VoIP service providers should be promoted, and mechanisms to enforce this cooperation should be investigated.

This list of inter-ITAD requirements is not limited to IP connectivity issues but covers all issues pertinent to offering global VoIP services. The AGAVE project will exploit these requirements and identify the ones related to the IP layer to be taken into account when designing VoIP-compliant Network Planes.

Note that the mechanisms to set up the cooperation between VoIP service providers are not the focus of the AGAVE project.

The term "Customer" within this section denotes VoIP service customer and the term "Provider" denotes a VoIP service provider. This latter can act under the "customer" or "provider" role with regards to SIAs. These two roles are not taken into account when identifying VoIP service provider requirements to build QoS-aware inter-ITAD VoIP service offerings.

5.2.1 Customer requirements

5.2.1.1 VoIP-C1: Global reachability

Customers should be able to reach every telephone number or VoIP destination independently of its location. By location, we denote the VoIP service provider the remote destination is attached to.

Specifically, local customers should be able to access a large set of destination numbers everywhere in the world, independently of the terminating VoIP service providers and whatever the scheme of the remote customer identifier is (e.g., E.164 number, H.323 URI, SIP URI, IAX URI, etc.). Moreover, the customers should be able to reach and to be reached from PSTN and other PLMN realms.

International Free Service and Telephony grades should be offered by VoIP Service Providers (see ITU recommendations for more information about International Free Services (IFS) and Telephony grades).

5.2.1.2 VoIP-C2: Transparency of inter-ITAD calls

Inter-ITAD calls should be transparent to all call participants. Inter-ITAD calls should be similar to intra-ITAD ones in terms of invocation procedures, perceived QoS, security level, anti-SPIT (Spam over IP Telephony) techniques like [SCHW06], etc. End users should not experience any difference from calls made from a party attached to the local VoIP service provider domain. The customer should not be aware that the call is traversing several telephony domains, independently of the number of ITADs crossed so as to reach the terminating domain.

5.2.1.3 VoIP-C3: Confidentiality and privacy

Customers require the confidentiality of their calls and the privacy of their identities. Service anonymity should be ensured and guaranteed. This anonymity should be provided in both IP and service layers. Media traffic should be encrypted and not transmitted as plain data. Signalling messages should also be protected. Customer terminals should be protected against Spam over IP Telephony and against (D)DoS attacks. Means should be deployed by VoIP service providers to detect identity usurpation and spoofing attacks.

The VoIP service provider should inform its customers on the information it retrieves from the user terminal equipment(s) for the needs of service delivery. Customers should own means to select their encryption options and associated security keys. Customers should be able to activate end-to-end security means for both media streams and signalling messages.

5.2.1.4 VoIP-C4: QoS

The perceived QoS, captured in Mean Opinion Score (MOS), of inbound and outbound calls should always be acceptable, irrespectively of the number of the crossed telephony domains. In other words, the involved IP and service layer mechanisms should ensure that the delay and loss introduced by crossing several BGP domains is negligible.

5.2.1.5 VoIP-C5: Availability

The availability is one of the key indicators of the robustness of the VoIP service. The VoIP service should be available in case of predictable system failures and especially for emergency calls in crisis situations and during disasters like earthquakes, etc. (see also next section). Service failures should be rare. VoIP service providers should deploy means to avoid "Avalanche Restart" and "Flash crowds" phenomena.

5.2.1.6 VoIP-C6: Emergency calls

Customers should be able to place emergency calls in emergency situations. Service providers should handle emergency signalling messages and associated media streams with high priority in both the IP network and service layers. Emergency calls associated data should be unambiguously identified so as to be protected and routed to the closest PSAP (Public Safety Answering Point). Customers should not have to provide location information. Instead customer location should be retrieved by the Service Provider itself (for instance deploying the DHCP (Dynamic Host Configuration Protocol) option 82 [PATR01]). Service providers should deploy means to check the validity and the consistency of location data and avoid the usage of static data for location information. For more information about

emergency calls, the reader is referred to ECRIT IETF Work Group documents, or to ITU Emergency Telecommunications Services related work.

5.2.1.7 VoIP-C7: Remote access

Customers should have access to the service even from locations remote to their home telephony domain. Two cases can be distinguished. The first case when the visited telephony domain is managed by the same VoIP service provider. And the second case, when the visited telephony domain is managed by another VoIP service provider. This second case is denoted as roaming. The customer profile and service features should be preserved. The perceived QoS of the calls should be similar to the one provided when using the home telephony domain. Inter-ITAD calls should be allowed in remote access and the support for emergency calls should be preserved.

5.2.1.8 VoIP-C8: Codec selection

Customers should be able to select the codec that they want to be used to send/receive their media streams (distinct codecs can be configured for receiving or sending only directions). This feature can be negotiated dynamically per call or statically upon subscription. This requirement is optional and depends on the capabilities of the nodes crossed by the media streams. In case of deployment of SBCs, the support of this requirement entails updating all these intermediate boxes so as to be compliant with most popular, efficient and robust codecs.

5.2.1.9 VoIP-C9: Local interface selection

End users should be able to determine the physical/logical interface which will be used to send or to receive VoIP traffic (both signalling and media streams). This requirement can be enforced either per call, per session, per registration session, etc. Consequently, the VoIP service provider should not bind statically the VoIP service to a physical/logical interface. This issue is essential when the customer is multi-homed or has several options of IP connectivity types (e.g., IPv4, IPv6, etc.) or QoS levels.

5.2.1.10 VoIP-C10: Heterogeneous access support

This requirement means that customers can access worldwide telephony services independently of the access technologies they are using and especially independently of the IP protocol version. Wifi, WiMax, ADSL, etc customers should have access to the service. Especially, no discriminations should be made between IPv4 and IPv6 customers. Heterogeneous session should be possible to be placed between IPv4 and IPv6 realms in the same conditions as for homogenous sessions, particularly the same level of QoS and security should be ensured.

5.2.1.11 VoIP-C11: Service assurance and monitoring

Customers should have the ability to verify the fulfilment of the SLA they subscribed to. Some indicators should be put at the disposal of the customers indicating the status of the service. The service assurance indicators and other measurement related information should be captured in the SLA. Billing tickets can be correlated to the indicator values.

Some of the indicators that may be provided to customers are listed hereafter:

- Availability of the service over the last period;
- Success rate of placed calls;
- Number of failures that happened over the last period;
- Loss and delay, etc.

In fact, service providers should provide information on their capabilities in terms of the same metrics for their service performance overall across all SLAs. This information should be available to customers before subscription to allow for the evaluation of the service provider.

5.2.2 Service provider requirements

5.2.2.1 VoIP-P1: Global coverage

A VoIP service provider is not able to provide global coverage of all telephony (VoIP) destinations. Therefore, VoIP service providers should extend their telephony service scope beyond the boundaries of their own domains (similar to the notion of international calls described in ITU E.105-110).

A VoIP service provider should be able to reach every telephone number or VoIP destination independently of its location. By location, we denote the VoIP service provider the remote destination is attached to. Specifically, VoIP service providers should maintain routes towards a large set of destination numbers everywhere in the world, independently of the terminating VoIP service providers and whatever the scheme of the remote customer identifier is (e.g. E.164 number, H.323 prefix, SIP URI, IAX URI, etc). Moreover, the VoIP service providers should be able to reach and to be reached from PSTN and other PLMN realms.

5.2.2.2 VoIP-P2: Support of numbering schemes other than E.164

VoIP service providers must be able, for inter-ITAD VoIP purposes, to support numbering schemes additional to E.164 and to construct and maintain routes towards these telephone numbers. Particularly, SIP URI, IAX URI, H.323 prefixes, e-mail addresses-like, and IP addresses should be supported. The logic and interpretation of these numbering schemes can be standardised or negotiated between interconnected VoIP service providers. Some of these number schemes can be exclusively associated with a given signalling protocol.

5.2.2.3 VoIP-P3: Discovery of VoIP providers and their capabilities

A VoIP service provider should be able to discover other VoIP service providers and their capabilities in order to establish SIAs. These capabilities may include a list of supported codec(s) whenever trans-coding is needed, capacity of remote service equipment in terms of number of calls supported simultaneously, the scope of IP connectivity, terminating telephone numbers, etc.

5.2.2.4 VoIP-P4: SIA flexibility

VoIP service providers should be provided with the means necessary to request, negotiate, establish, modify and delete SIA agreements. The establishment of an SIA should not freeze the ability of the provider to release, modify and delete previously existing SIAs. The modification of a given agreement should not disturb other agreements relying on resources provided by the agreement under modification.

5.2.2.5 VoIP-P5: Interoperability

Adjacent VoIP service providers are not assumed to deploy the same signalling protocol(s). Thus, signalling gateways, to translate the incoming/outgoing protocol messages, should be supported by one of the two interconnected service peers. In addition, trans-coding functions should be implemented. The logic of invoking these translation/trans-coding functions can be configured statically or dynamically. The two service peers should agree on how these functions are to be invoked.

5.2.2.6 VoIP-P6: Exchange of homogenous call routing data

In order to offer universal VoIP services, interconnected VoIP service providers must exchange routing data so as to be able to forward signalling messages. Routing data should be updated and validated, thus it is advisable to use automated dynamic routing mechanisms instead of configuring static routes. As an automated and dynamic process, call routing data exchange eases the discovery and advertisement of remote/local telephony destinations with additional information on metrics relevant to the route selection process.

Only authorised service providers should have access to the routing data. Only two service peers would share routing data in a bi-lateral SIA or all members of a federation can share their routing information. A description of federation based architecture is provided in [HABE06]. Authentication, authorisation and integrity mechanisms should be deployed so as to accept routing data as valid. A VoIP service provider may advertise its own or its peers' telephony prefixes.

5.2.2.7 VoIP-P7: Ability to tune the call route selection process

Each VoIP service provider should be able to tune and configure its own call route selection process and to tune its advertisement policies (export/import policies).

5.2.2.8 VoIP-P8: Support of multiple call routing paths

In order to load balance its traffic or to enforce some service engineering operations, each VoIP service provider should be able to receive/advertise multiple call routing paths. As stated above, these paths may be used to load balance the telephony traffic, or to backup the primary path in case of congestion or failures. Alternative paths for the same type of traffic must be equivalent in terms of QoS treatment for the media flows. This is similar to multi-homing scenario in the IP world.

5.2.2.9 VoIP-P9: Optimisation of signalling and media paths

In order to reduce the set-up time of inter-ITAD calls, VoIP service providers should be able to select the optimum end-to-end signalling path. This means, that the number of crossed ITADs should be minimal. The aim is to reduce the delay that may be introduced due to specific (service) operations in VoIP nodes residing in crossed VoIP domains, and also to minimise failures (more precisely probability of failures). Another sub-requirement is to reduce the amount of inter-ITAD signalling control messages exchanged between interconnected VoIP domains in order to place/teardown an inter-provider call.

In addition, the end-to-end IP path should be optimum and within the constraint that the end-to-end perceived QoS meets the customer requirements.

5.2.2.10 VoIP-P10: Resilience and availability

In order to implement resilient and close to 5 9's available inter-ITAD VoIP service offerings, VoIP service providers should:

- Deploy means to detect and repair failures and maintain a minimal service level in both service and IP connectivity layers;
- Avoid single points of failure, deploy redundant interconnection equipment;
- Ensure end-to-end resilience;
- Be able to evaluate availability of the VoIP service end-to-end.

5.2.2.11 VoIP-P11: Synchronise service layer and control layer

Synchronising the service layer and the control layer is one of the critical issues related to the deployment of QoS-enabled VoIP service offerings. Indeed, service routing logic and underlying IP connectivity tuning logic should be consistent so as to meet the requested QoS treatment. For more information about this synchronisation issues, the reader is referred to section 4.3.2.7.2.

5.2.2.12 VoIP-P12: Ability to detect INP spirals

By INP spiral, we mean that an inter-provider (inter-ITAD) call crosses an INP domain several times. VoIP service providers should have means to avoid this phenomenon or at least to be aware that it will occur when selecting a given inter-ITAD path. For more information about INP spiral, please refer to section 4.3.2.7.1.

5.2.2.13 VoIP-P13: Ability to evaluate the QoS along an inter-ITAD path

Before selecting a given inter-ITAD path, a VoIP service provider should be able to evaluate the QoS level that the call media streams will experience over this path. Valid QoS-related information should be at the disposal of the VoIP service provider so as to drive the selection of the optimum next VoIP service provider to be contacted. This QoS information should reflect the real or closest to the real status of the service QoS level and not remain static.

A VoIP service provider may be informed if congestion is building up along the selected path. Appropriate means to avoid congestion are to be investigated.

5.2.2.14 VoIP-P14: O&M

VoIP service providers should be able to run FCAPS functions for inter-ITAD calls. The required enhancements should be manageable and not complex to implement. Service peers should agree if exchange of monitoring/operational/accounting data is required. Each service provider may deploy its own O&M architecture. A common O&M architecture across all domains is not required.

Means to detect failures at inter-ITAD links should be provided by the underlying INPs, or deployed between the interconnected service peers.

5.2.2.15 VoIP-P15: Billing for inter-domain calls

Service providers should agree on how charging, billing and accounting are to be handled between their domains. Financial models, such as transit or peering, need to be agreed between service peers.

5.2.2.16 VoIP-P16: SIA assurance and monitoring

A VoIP service provider should have means to monitor the usage of each SIA and whether a service peer meets its contractual commitments. VoIP service providers should have the ability to check the fulfilment of their obligations (role of client) and evaluate if the service peers (role of provider) meet theirs as agreed in the SIA. Valid indicators and data should be put at the disposal of the (client) VoIP service providers to verify the level of the established interconnection. This data should be updated to reflect the real status of the interconnection and it must be described and listed in the SIA. Billing tickets can be correlated to the indicators values.

In addition, the exchange of monitoring data, monitoring methodology, monitoring templates, etc should be agreed between the two service peers.

5.2.2.17 VoIP-P17: Support of "Import" and "Export" policies

By "Import" and "Export" we denote the ability of a VoIP service provider to indicate to or to request from a service peer the policies to apply for a particular SIA. For instance, a VoIP service provider can indicate to its service peers not to advertise to him paths including a given ITAD identifier, or only to advertise paths/routes crossing two telephony domains. Additional policies can be agreed and configured by both interconnected service peers. Additional information exchange for inbound and outbound service engineering may be investigated.

5.2.2.18 VoIP-P18: Security

The inter-ITAD architecture must support both hop-by-hop and end-to-end security modes as the ones described, for instance, in [ROSE02]. In addition, the inter-ITAD solution may support end-to-middle security schemes as described, for instance, in [ONOK05a] and [ONOK05b]. Both signalling and media streams must be secured. At least SRTP [BAUG04] and IPSec [KENT98a] [KENT98b] should be considered as candidates to secure media streams. Additional features like anti-spoofing, anti-registration/subscription hijacking attack detection and means against usurpation should be supported by the inter-ITAD architecture.

In addition, in case of interworking between heterogeneous realms (IPv4, IPv6, etc.), security mechanisms should not be broken.

5.2.2.19 VoIP-P19: Ensure private communication between service nodes

A best current practice of VoIP service providers is the configuration of VPNs to interconnect VoIP service nodes, as means to protect service nodes and to isolate the VoIP signalling traffic. This intra-domain practice may be extended to inter-ITAD context.

5.2.2.20 VoIP-P20: Support of privacy and confidentiality

Confidentiality and privacy must be preserved when crossing several telephony domains. Therefore, VoIP service providers should ensure with their service peers that dedicated means have been deployed so as to protect the confidentiality of inter-domain calls and/or associated media streams. The level of ensured confidentiality and privacy must be at least the same as that guaranteed for intra-ITAD calls. In addition, calls anonymity should be guaranteed in both service and IP layers. At the IP layer, the IP addresses of the callee should not be "seen" by the caller and vice-versa unless this is explicitly requested by the customers.

5.3 IP/MPLS-based VPNs

The purpose of this section is to identify the VPN-specific requirements in order to offer VPN services with world-wide coverage. For more information about the VPN use case, please refer to section 4.3.3. Note that an initial set of inter-AS VPN requirements has been identified and elaborated within the context of MAVS initiated and captured in [HALS06] Internet-Draft. The requirements described below do not rely on specific protocol to interconnect VPN service providers.

Within the context of this deliverable, we assume that no single VPN service provider can ensure global reachability. From this standpoint, cooperation between VPN service providers should be promoted, and mechanisms to enforce this interconnection should be investigated.

5.3.1 Customer requirements

5.3.1.1 VPN-C1: QoS transparency/translation at the customer level

There is general agreement that customer packets should not be remarked (that is, have their DSCP values modified) as they cross the VPN service provider (VPN SP) domain. At the same time, it is often necessary for the VPN SP to impose a QoS treatment on customer packets that differs from that which might be indicated by the customer's DSCP (such as is the case for non-conforming traffic downgraded to best effort). However, even if the packets are treated as best effort ones by the VPN SP, the VPN customer wishes to retain the original DSCP marking for its own uses when the packets arrive at the remote site(s). This requirement is defined as "QoS transparency" (a.k.a. DSCP transparency), referring to the tunnelling of the customer QoS policy signalled by the DSCP, unmodified, from ingress site to egress site across the VPN SP domain.

Two VPN sites of the same customer usually have common QoS policies based on a common interpretation/enforcement of packet marking (mainly the IP ToS field). The VPN SP should not use this field for its own QoS policies as the VPN customer already uses it and does not want it to be modified.

In the case of an extranet when one site of a VPN communicates with a site of another VPN, for a given QoS level, the two VPNs usually have a different usage of the IP ToS field. The customers may then require the re-marking of the ToS field to be enforced as an additional service by the VPN SP so as to ensure the consistency of the QoS treatment between the two VPN sites.

5.3.1.2 VPN-C2: VPN topology

The VPN customer should be able to control the topology of its VPN(s). Usual topologies are:

- Full mesh or any to any, where each site can directly communicate with all other sites;
- Partial mesh;
- Hub and spoke, where the communication between two spoke sites should go through a hub site, for example for security reasons;
- Extranet where a site of a VPN needs to communicate with a site of another VPN.

These topologies can actually be different from the topology of the VPN SP and they should not be constrained by the geographic scope of the VPN SP.

5.3.1.3 VPN-C3: Internet access

From its VPN SP, a customer may ask additionally Internet access, usually with added value services such as NAT traversal and firewall rules to protect its VPN sites from connections coming from outside the VPN (i.e. the Internet).

5.3.1.4 VPN-C4: Global reachability

Customers should be able to attach any site to a specific VPN independently of the location of the site. By location we denote the VPN SP a customer site is attached to.

The customer's requirement is access to a VPN to be as ubiquitous as access to the Internet.

5.3.1.5 VPN-C5: VPN access means

The customers need to be provided with appropriate means to initiate and terminate access to a specific VPN, within the invocation and authentication profiles agreed between the VPN customer and the VPN SP at subscription time.

This requirement entails the following:

- Customers should be able to get access to a VPN either explicitly or implicitly. Explicit invocation will yield the use of an explicit authentication protocol. Implicit invocation is applicable to permanently activated services and does not require the explicit use of an authentication protocol and usually relies on trusting the layer 2 access network (such as the physical access line for fixed access or the GSM/SIM authentication for mobile access).
- Customers should be provided with appropriate invocation and authentication means to access a specific VPN.

From SP VPN perspective, this requirement yields the following needs:

- Providers should provide for automated means for authenticating and authorising an implicit or explicit request for accessing a VPN service.

5.3.1.6 VPN-C6: Self-monitoring means

Customers should be able to check the quality of the VPN services, in order to verify among others that there is no intrusion, i.e. the sites/hosts granted access to the VPN were all authorised. This requires that they should be provided with appropriate self-monitoring means.

From a provider's perspective, this requirement yields the following requirements:

- Providers should provide customers with appropriate means, enabling the customers to inspect the quality of the service and the VPN users with access to the service, active users or historical records.
- Providers should cater for appropriate customer care means for receiving and analysing customer inquires and complaints with respect to the service compliance to the agreed terms.

5.3.1.7 VPN-C7: Transparency to inter-SP VPN

The traffic should cross multiple VPN service provider domains transparently for all VPN sites. Inter-SP flows should be treated similarly (in terms of invocation procedures, perceived QoS, security level) to intra-SP ones. The complexity of establishing inter-SP VPNs should be handled at the service provider side. End users should not experience any difference from traffic coming from a party attached to the local SP domain. The customer should not be aware that the traffic is traversing several VPN SP domains and/or INP domains.

5.3.1.8 VPN-C8: Availability

The availability of the VPN service is a key requirement from VPN customers as VPN carries data of mission critical applications. Key indicators are global availability and the Mean Time To Repair (MTTR) in case of failure.

5.3.1.9 VPN-C9: Service assurance and monitoring

A VPN customer should have means to monitor the service and evaluate whether the service provider meets the terms agreed in the SLA.

An SLA may be defined per access network connection, per VPN, per VPN site, and/or per VPN route. In some cases the guaranteed levels for SLA parameters may depend upon the scope of the VPN. For example, one level of guarantees might be provided for service within a single VPN SP, whereas a different (generally less stringent) level of guarantees might be provided for VPNs crossing multiple SPs.

5.3.1.10 VPN-C10: Management

A customer should have means to view the topology, operational state, service status, and other parameters associated with its VPNs. The customer should be able to configure and tune all aspects of management information about CE devices and VPN customer-specific attributes managed by the VPN SP.

5.3.1.11 VPN-C11: Load balancing

To optimise resource utilization between their sites, customers should have load balancing means on inbound and outbound traffic to/from site dual attached to a VPN SP or dual homed to two VPN SPs.

5.3.2 Service Provider requirements

5.3.2.1 VPN-P1: QoS transparency/translation at the SP level

Two network nodes of the same VPN SP usually have common QoS policies based on a common interpretation/enforcement of packet marking (mainly the IP ToS field) even when they are not adjacent but interconnected through several INPs. The VPN SP marking should be carried transparently across the several INPs. In the case of an inter-SP VPN, for a given QoS level, the two VPN SPs usually have a different usage of the IP ToS/MPLS EXP field. A mechanism should be defined to ensure QoS consistency and translation of the IP ToS/MPLS EXP fields when two or several VPN SPs are involved in order to offer a single VPN.

5.3.2.2 VPN-P2: Internet access inter-SP optimisation

In order to reduce transit delay and reduce the cost of using an expensive QoS-aware SP/INP infrastructure for best effort Internet traffic, when a VPN spans across multiple SPs and the customer needs Internet access, Internet flows may be routed to the Internet to the nearest Internet peering of the first SP receiving the traffic.

5.3.2.3 VPN-P3: Global coverage

A local VPN service provider should be able to provide global coverage in terms of location of the customer sites. Therefore, VPN service providers should extend their VPN service scope beyond the boundaries of the INP(s) where their points of presence are located.

5.3.2.4 VPN-P4: Discovery of VPN providers and their capabilities

A VPN service provider should be able to discover other VPN service providers and their capabilities in order to establish SIAs. These capabilities may include geographic coverage, QoS capabilities, supported tunnelling technologies, added value services like Internet access, NAT, firewall, etc.

5.3.2.5 VPN-P5: SIA flexibility

VPN service providers should be provided with the means necessary to request, negotiate, establish, modify and delete SIA agreements. The establishment of an SIA should not freeze the ability of the VPN SP to release, modify and delete previously existing SIAs. The modification of a given agreement should not disturb other agreements relying on resources provided by the agreement under modification.

5.3.2.6 VPN-P6: Interoperability

Adjacent VPN SPs are not assumed to deploy the same tunnelling protocols or VPN QoS. Thus, VPN ASBRs able to translate the incoming/outgoing packets, should be supported by the service peers. The logic of invoking these translation functions can be configured statically but preferably dynamically. The two service peers should agree on how these functions are to be invoked.

5.3.2.7 VPN-P7: Support of multiple SP paths

In order to load balance its traffic or to enforce some service engineering operations, each VPN service provider should be able to receive/advertise multiple VPN paths. These paths may be used to load balance the traffic, or to backup the primary path in case of congestion or failures. Alternative paths for the same type of traffic must be equivalent in terms of QoS treatment.

5.3.2.8 VPN-P8: Resilience and availability

In order to implement resilient and close to 5 9's available VPN service offerings, VPN service providers should:

- Deploy means to detect and repair failures and maintain a minimal service level at both service and IP connectivity layers.
- Avoid single points of failure, deploy redundant interconnection equipment;
- Ensure end-to-end resilience;
- Be able to evaluate availability of the VPN service end-to-end.

5.3.2.9 VPN-P9: O&M

VPN SPs should be able to run FCAPS functions for inter-SP traffic. The required enhancements should be manageable and not complex to implement. Service peers should agree if exchange of monitoring/operational/accounting data is required. Each service provider may deploy its own O&M architecture inside its domain.

5.3.2.10 VPN-P10: Billing for inter-domain traffic

VPN SP should agree on how charging, billing and accounting are to be handled between their domains. Financial models, such as transit or peering, need to be agreed between service peers.

5.3.2.11 VPN-P11: SIA assurance and monitoring

A VPN service provider should have means to monitor the usage of each SIA and whether a service peer meets its contractual commitments. VPN SPs should have the ability to check the fulfilment of their obligations (role of client) and evaluate if the service peers (role of provider) meet theirs as agreed in the SIA. Valid indicators and data should be put at the disposal of (client) VPN SP to verify the level of the established interconnection. This data should be updated to reflect the real status of the interconnection and it must be described and listed in the SIA. Billing tickets can be correlated to the indicators values.

In addition, the exchange of monitoring data, monitoring methodology, monitoring templates, etc should be agreed between the two service peers.

5.3.2.12 VPN-P12: Security

The inter-SP architecture must support security means as this is the main service of a VPN (beside isolation).

At minimum, the architecture must support isolation between VPNs at the control and forwarding planes. It must also support isolation between a VPN and other SP/INP flows and in particular from the Internet. An SP should be able to control the list of SPs who participate to a VPN and excluding other SPs from gaining access to that VPN.

The architecture should support for authentication between the customer and the SP, between SPs, and between SP and INP. Additionally, it may support encryption of VPN control and/or traffic flows.

5.3.2.13 VPN-P13: Scalability

Scalability is a key concern for a telecom provider whose financial goal is to profit from its expensive infrastructure by increasing the number of customers.

On one hand existing intra-SP VPN services already face scalability issues especially regarding the number of VPN routes handled by the SP. On the other hand, the existing worldwide Internet has also scalability issues regarding the number of routes and their instability. The combination for providing worldwide inter-SP VPN services, will certainly raise scalability issues.

Key scalability parameters are the number of VPN SPs, VPN nodes, customer sites, VPN routes and VPN routing instability.

A mechanism should allow to filter or aggregate information between SPs.

5.3.2.14 VPN-P14: Stability

In addition to availability and scalability, a VPN service should also be stable.

In addition to tunnel stability, stability is a property of several functions such as VPN routing, signalling and discovery mechanisms. For example, in the case of routing, route flapping or routing loops should be avoided in order to ensure stability. Stability of the VPN service is directly related to the stability of the mechanisms and protocols used to establish the service. It should also be possible to allow network upgrades and maintenance procedures without impacting the VPN service.

5.3.2.15 VPN-P15: Resource sharing

Network resources such as memory space, FIB table, bandwidth and CPU processing should be shared between VPNs and, where applicable, with non-VPN Internet traffic. Mechanisms should be provided to prevent any specific VPN from taking up available network resources and causing others to fail. Guarantees to this effect should be provided to the customer in the SLAs.

Similarly, resources used for control plane mechanisms are also shared. When the control plane distributes VPN specific information and provides other VPN control mechanisms, there shall be mechanisms to ensure that control plane performance is not degraded below acceptable limits when

scaling the VPN service, or during network events such as failure, routing instabilities etc. Since a service provider's network may be used to provide Internet service, in addition to VPNs, mechanisms to ensure the stable operation of Internet services and other VPNs shall be made in order to avoid adverse effects of resource hogging by large VPN customers.

5.3.2.16 VPN-P16: Management

Service provider should have means to view the topology, operational state, service status, and other parameters associated with each customer's VPN. Furthermore, the service provider should have means to view the underlying logical and physical topology, operational state, provisioning status, and other parameters associated with the equipment providing the VPN service(s) to its customers.

In the multi-provider scenario, it is unlikely that participating providers would provide each other a view to their network topology and other parameters mentioned above. However, each provider should ensure via management of their own networks that the overall VPN service offered to the customers is properly managed. In general the support of a single VPN spanning multiple SP domains requires close cooperation between the service providers. One aspect of this cooperation involves agreement on what information about the VPN will be visible across providers, and what network management protocols will be used between providers.

6 AGAVE HIGH-LEVEL SPECIFICATIONS

This section provides high level specification of Network Planes and Parallel Internets. It describes also interactions between service actors so as to build IP (QoS-enabled) services across several Network and Service Providers.

6.1 Interactions between Service Actors

Customers, SPs and INPs establish vertical (SLAs and CPAs) and horizontal (SIAs and NIAs) bilateral agreements (see sections 3.1.2, 3.1.3 and 3.1.4) in the context of which they interact to achieve their business objectives.

Agreements may follow either the customer-provider paradigm or the peer partnership paradigm. Vertical relationships, i.e., customer to SP and SP to INP, follow the customer/provider paradigm. Horizontal relationships may be of either type, depending on the hierarchical level of the actors, e.g., a tier-2 INP is the customer of a tier-1 INP while it is a peer partner of another tier-2 INP.

The cycle of interactions between actors spans through the following phases:

- *advertising and discovery* phase where each actor advertises its capabilities and seeks other actors to establish agreements with
- *negotiation and agreement establishment* phase where two actors instantiate and negotiate potential agreements and eventually undertake the necessary actions to establish and activate the reached agreement
- *operation* phase where the involved parties interact following the terms and conditions contained in the established agreement; during operation each party may receive feedback information on the behaviour of the service/interconnection and may initiate tuning actions to be carried out immediately, or issue invocation and agreement modification requests subject to negotiation
- *assurance* phase where the degree to which each party comply with the terms of the agreement is evaluated, following appropriate procedures as part of the agreement

Service or interconnection agreements must hence contain information on the service itself, the expected feedback and the permissible actions during operation, as well as information required for the activation and the assurance phases.

The following sections elaborate on the interactions between the service actors.

6.1.1 INP to INP

For the purpose of expanding the scope of their IP connectivity, INPs interact with each other on the basis of *INP Interconnection Agreements* (NIAs) (see section 3.1.2).

A NIA can be established between adjacent or remote INPs (see Figure 20). The latter is called a *remote NIA*. Remote are those INPs with no adjacency to the reference INP. A remote NIA is required when adjacent INPs do not provide the desired level of QoS all the way down to remote destinations.

An INP may provide guarantees for reaching directly attached sites only or for remote sites too. Remote are those sites not directly connected to the reference INP. The latter is then called a *multi-hop NIA* as the INP provides guarantees not only to local destinations, but also to destinations multiple AS hops away. The remote sites specified in a multi-hop NIA may extend as far as the end destination sites.

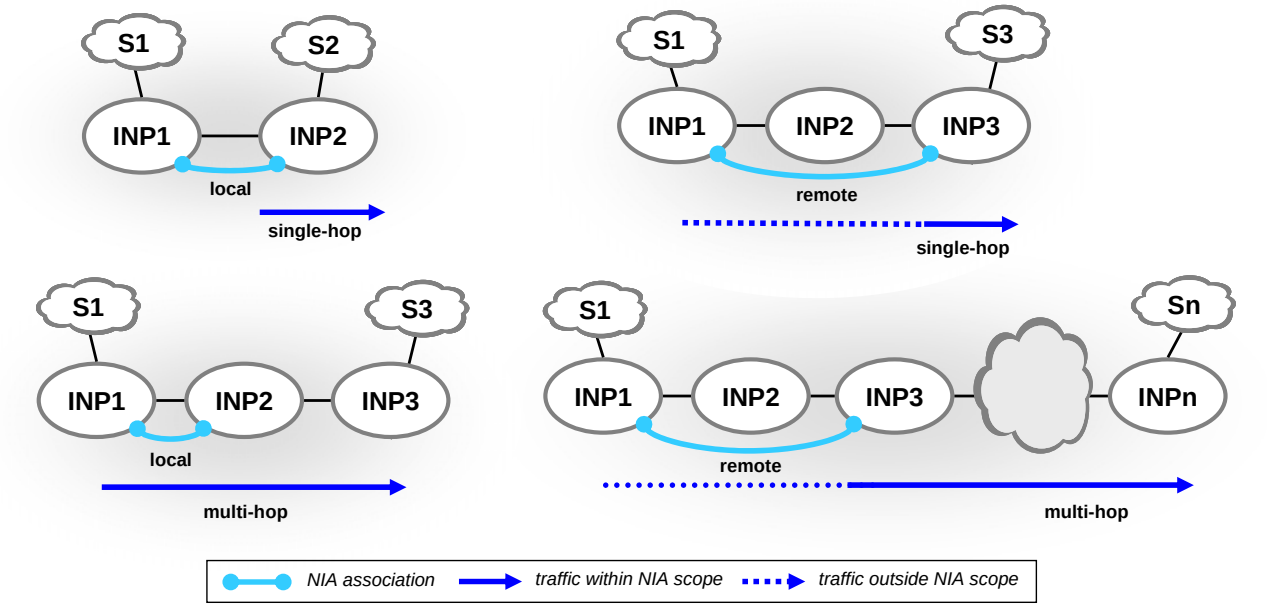

Figure 20 INP interconnection agreement

The association between two INPs may follow the customer-provider or the peer partnership paradigm, mostly depending on the size of the INP (see discussion in section 3.1.2). Under the customer-provider paradigm the customer INP buys connectivity to destinations reached by the provider INP and pays the cost of the links to the provider INP. In the peer partnership paradigm the associated INPs mutually exchange connectivity to the destinations reached by each, sharing the cost of their physical interconnection.

6.1.1.1 NIA Content

An NIA is drafted and further elaborated during negotiation by the interacting INPs. The information contained in the NIA includes:

- *Interconnection Connectivity* information specified in terms of:
  - *Interconnection Scope*: sets the boundaries across which the NIA guarantees apply, expressed in traffic ingress and egress points. The egress points are local for single-hop and remote for multi-hop NIAs. An ingress/egress point can be specified in terms of ASBR interface(s), as a set of IP address prefixes identifying remote hosts or sub-networks, or using descriptive labels that can be deterministically translated to IP address prefixes such as geographical areas.
  - *Ingress Flow Identification*: provides the classification rules identifying the traffic to receive the treatment specified in the NIA. Classification is performed based on the IP header fields (e.g., source and/or destination IP address, protocol, ToS/DSCP, etc.), and/or based on tunnel end identifier if tunnelling is used for interconnection. Multiple ingress flow identifiers are used in the case of providing different levels of QoS and resilience guarantees, for signalling which specific level of guarantees the traffic is to get. The specific identifiers are determined by the downstream INP controlling the available identifiers on its interfaces.
  - *Egress Flow Identification*: sets the marking and/or tunnelling rules to be applied to the traffic at the egress of the NIA. It is required when another remote NIA is to follow in the downstream path; the egress flow identification is used as the ingress flow identification for the INP following the egress point of the reference NIA.
  - *Flow Identification Map*: specifies the association between ingress and egress flow identifiers. An ingress identifier must be associated with exactly one egress identifier and many ingress identifiers can be associated with the same egress identifier. A distinct ingress-egress flow

identification pair is introduced for each overlaying NIA in case of identification delegation (see section 6.1.1.5.1).

- *Incoming Traffic Conformance*: specifies the traffic envelope to which the traffic identified by ingress flow identification should comply to in order to receive the agreed guarantees.
- *Outgoing Traffic Conformance*: specifies the traffic envelope to be enforced by policing or shaping the traffic leaving the INP's network at an egress point. Outgoing traffic conformance is offered as an optional feature.
- *Performance Guarantees* and *Excess Treatment*: specify the treatment the traffic will receive as a result of the NIA. Performance guarantees are provided for the conformant traffic while excess treatment is applied to non conformant traffic. Performance guarantees are specified in terms of packet transfer performance metrics (delay, loss and jitter) and throughput. Performance guarantees are expressed either quantitatively for statistical and hard guarantees or qualitatively, e.g., better-than-best-effort, silver, gold. Excess treatment can be shaping, dropping or degraded packet transfer performance.
- *Resilience Guarantees*: specify the guarantees for the agreed performance to be provided reliably, in terms of acceptable failure frequency and repair time per type of eventuality, e.g., for single link failures.
- *Security Guarantees*: specify the security requirements for carrying traffic across the scope of the NIA in terms of data confidentiality, integrity and replay attack prevention.

- *Activation* information for interactions that need to take place between the INPs to complete the NIA activation beyond internal INP configuration, either between adjacent or remote INPs. Such interactions may be establishing a peering connection between BGP speakers or between invocation protocol speakers for dynamic tunnel establishment and teardown, establishing of security associations between border routers at the NIA interconnection points.
- *Operational Feedback* information. To perform fault and performance management each INP requires feedback from the network, including the interconnections with other INPs. The particular requirements are captured in the following generic clauses:
  - *Monitoring Tasks*: sets the monitoring reports and/or alarms that the downstream INP must produce and deliver to the upstream INP periodically or on-demand. Applicable reports/alarms are related to network performance, failure incidents, security attacks, troubleshooting logs etc. (see section 6.3 for details). A monitoring task is specified in terms of the metrics (e.g., loss, delay), data collection features ,e.g., granularity, sampling frequency, and finally, data dissemination features e.g. reporting frequency, alarm thresholds etc. The scope of a monitoring task may be limited to ingress INP intra-domain or it may include inter-domain statistics extending as far as the final destinations or the egress INP for multi-hop NIAs.
  - *Probing Facilities*: sets the probing scope within the interconnection scope, the terms and conditions for initiation and termination of active monitoring jobs spanning the probing scope.
- *Permissible actions* information. Typical examples of NIA modification are the expansion of the interconnection scope, either to local or remote interconnection points, and the increase of the reserved capacity. Applicable invocation requests include dynamic tunnel establishment and teardown, or bandwidth allocation and release in the case of a managed bandwidth NIA. Information required to specify potential NIA tuning, invocation and modification procedures includes:
  - *Request Means and User Info*: specify the procedures, the protocols and the authentication information for the users entitled to issue a tuning action, to invoke or to modify a particular instance or aspect of the NIA.
  - *Response Info:* specifies the maximum response time to and the probability for getting through invocation or modification requests of a particular instance or aspect of the NIA.

- *Assurance* information to specify the procedures available and the applicable parameters for assessing compliance of each party to the agreed terms and conditions.

6.1.1.2 Advertisement and Discovery

INPs advertise their capabilities in the context of potential NIA offerings in terms of:

- reachable local and remote IP address prefixes;
- QoS, resilience and security guarantees per group of IP address prefixes;
- tunnelling technologies at INP's edges (e.g., MPLS, IP tunnelling);
- support of ToS/DSCP marking and classification at INP's edges;
- security mechanisms at INP's edges (e.g., IPSec, encryption algorithms, etc.);
- operational feedback facilities (metrics, granularity, polling frequency, scope);
- set of NIA permissible actions;
- assurance procedures.

Discovery of the advertised capabilities of other INPs may be based on directory listings or on a protocol for bilateral exchange of NIA-related capability advertisements.

6.1.1.3 NIA Negotiation

Reasoning on a proposed NIA involves validating the NIA against static INP capabilities derived from equipment limitations such as support of a particular technology as MPLS, and for the NIAs found valid assessing if current equipment configuration and network dimensioning can accommodate the requested capacity, QoS and resilience guarantees, or if and which adjustments on network dimensioning could. In the latter case, internal and/or interconnection with peer INPs re-dimensioning may be triggered to accommodate the new request.

Re-dimensioning may involve re-shuffling capacity among existing network planes and paths, or even introducing new network planes, internally or further downstream with peer INPs for accommodating the requested capacity and guarantees. The extent to which resource dimensioning is adjusted on the received NIA requests depends on the tools and the policies of each INP.

6.1.1.4 NIA Activation

Activation of an NIA from the downstream provider entails the following actions:

- possible internal and/or interconnection re-dimensioning (see section 6.1.1.4 above);
- establishment of routing entries/tunnels specific to the NIA to ensure the traffic of the NIA will exit through the NIA egress points for the specified destination IP address prefixes;
- set-up of the NIA egress traffic marking and policing;
- set-up of the NIA ingress traffic classification and policing;
- configuration of internal mechanisms to provide operational feedback, to handle permissible actions and to activate assurance procedures;
- following of agreed activation procedures, e.g., establishment of security associations at interconnection points.

6.1.1.5 INP Interconnection Issues

6.1.1.5.1 Identification Delegation in Multi-Hop NIAs

Traffic identification on the ingress of a remote NIA following a multi-hop NIA must be enforced by an INP which has no NIA with the source INP (see Figure 21) unless the head INP uses direct tunnelling to the egress INP. In fact, all intermediate INPs must generate a new identifier for the new macro-flow except those transparently crossed over tunnels.

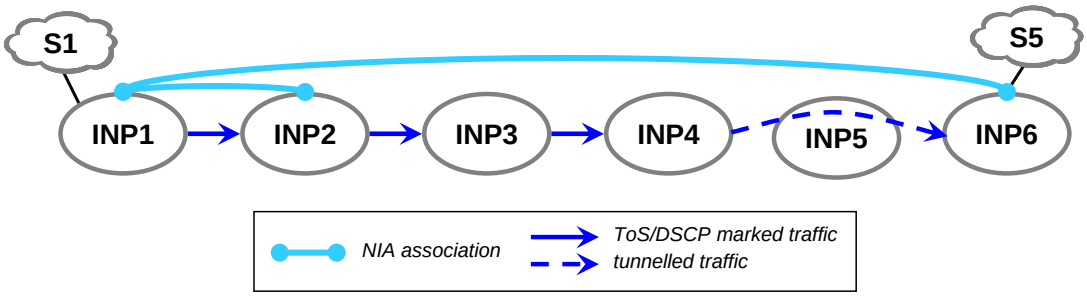

Figure 21 Identification Delegation in Multi-Hop NIAs

In the above example, identification on the ingress of the remote NIA (NIA between INP1 and INP6) following the multi-hop NIA (NIA between INP1 and INP2 covering INP2 to INP5 hops) must be enforced by an INP (INP5 or INP4 in the example as tunnelling is used) which has no NIA with the source INP (INP1). So, INP4 enforces the identifier required at the ingress of INP6 to identify traffic coming from INP1. INP2 being the head of the multi-hop NIA delegates traffic identification to intermediate INPs which all, except transparently crossed INP5, need to establish new ToS/DSCP or tunnel end identifiers to identify the traffic coming from INP1. Note that traffic between INP2 and INP6 is carried over existing NIAs among the INPs in the path; the only change is that part of the traffic will enter and exit each INP using an additional ingress/egress identification pair.

There are considerable scalability issues with egress identification delegation because of the interactions required between all intermediate INPs to establish additional identifiers and of the limited number of available identifiers. Note though that the identification delegation scenario is not expected to be commonly encountered and that the number of INPs in a path is expected to be small (the current average AS path length is 3.5 [RIPE06]).

6.1.1.5.2 Requirement for Differentiated Routing Support

When the NIA scope does not reach the final destinations, i.e., when the egress point is not the interface where the destination site is attached, then a path is required to direct the traffic for the NIA heading to the destination through the egress point. In case when INP internal routing policies advocate another route for the destination, differentiated routing mechanisms are required to maintain multiple routes to the destination. In the example of Figure 22, INP2 maintains a route for the traffic coming from INP1 to site S5 through INP3. Another route to S5 may be established for local traffic through INP4.

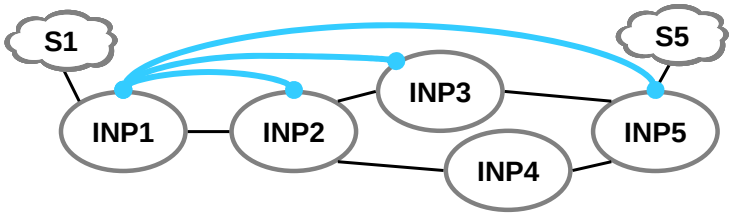

Figure 22 Differentiated Routing Requirement

In addition to the route for the local traffic, a distinct route is required for each NIA with traffic heading to this destination. In the absence of differentiated routing enabling mechanism, no NIAs with egress point before the final destination can be established, unless there is a single path to the

destination used inevitably by both the local and the transit NIAs traffic. Differentiated routing can be implemented either using tunnelling techniques or by one of the emerged multi-topology routing protocols.

6.1.1.5.3 Providing Path Control

An INP may have multiple alternative paths to reach some destinations, either internal and/or external paths. It may be useful to gain control in load distribution across alternative concurrent paths in case of NIAs of loose, qualitative QoS or resilience guarantees, as the QoS experienced across different paths may vary arbitrarily.

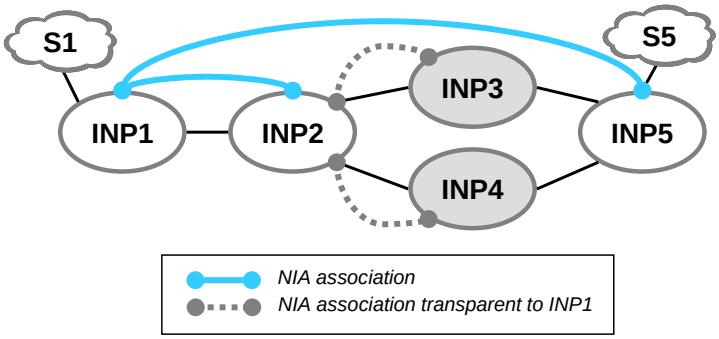

Figure 23 Providing Path Control

In Figure 23, INP2 has two alternative paths to site S5 and may be interested in providing path control to its upstream INP1. A unique identifier must be associated to each path at INP2's ingress and INP1 will control which packets are sent through which path by using the corresponding identifiers. As part of operational feedback interactions INP2 could provide performance data per path for INP1 to use for tuning load distribution.

Providing path control is not considered further as no concrete scenario arises for its use by the identified AGAVE service use cases.

6.1.2 SP to INP

SPs rely on the connectivity offered by INPs for fulfilling the connectivity aspects of their services. Beyond the forwarding and QoS treatment of the data traffic entering the INP's network from the SP's sites, the INP offers to the SP means to control the connectivity provisioning. The *Connectivity Provisioning Agreement* (CPA) between SP and INP regulates the relevant issues under the customer-provider paradigm (see section 3.1.3).

Depending on the service and the business policies of the SP, its connectivity provisioning requirements may significantly vary. A pipe interconnecting two remote SP sites completely controlled and guaranteed by the INP is at one extreme. The opposite extreme is an SP with the requirement to operate and manage a chunk of the INP's network, to define the topology, routing and forwarding over this network portion, to use the network control and management functions deployed by the INP for this network portion.

Equivalently to the INP Interconnection Agreement (see section 6.1.1), a CPA can be local or remote and single or multi-hop (see Figure 24). Note that the interconnected sites may be under the administration of the SP or they may belong to other SPs.

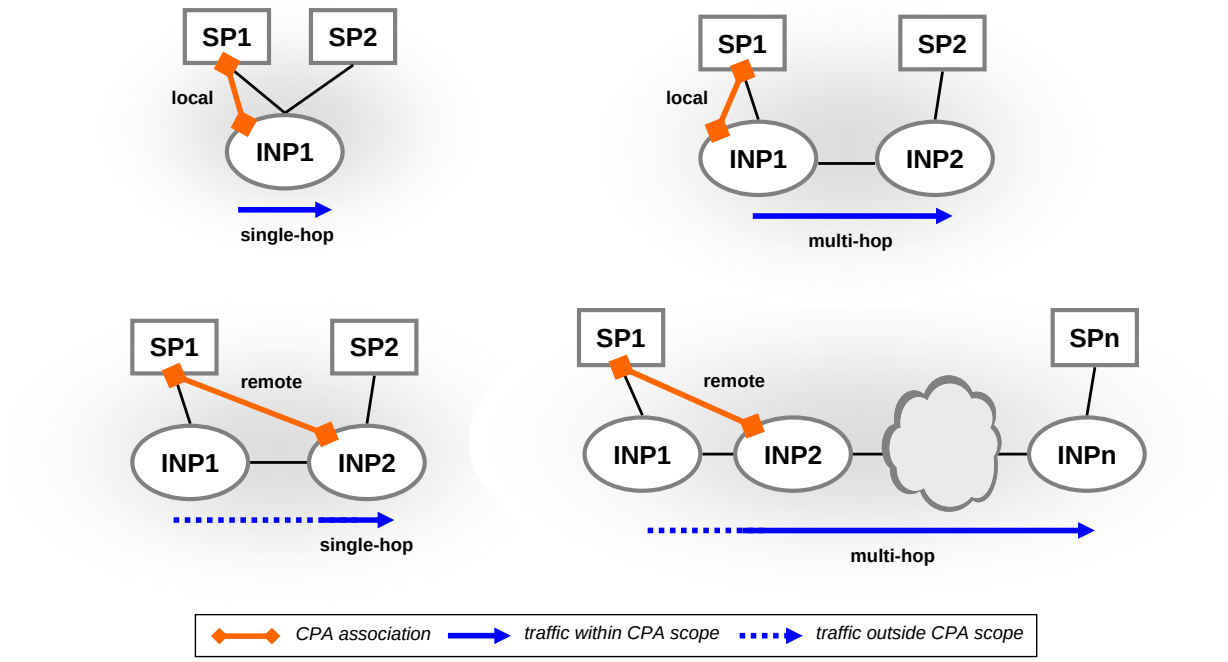

Figure 24 Connectivity provisioning agreement

6.1.2.1 CPA Content

The information contained in the NIA (see section 6.1.1.1) for capturing the interconnection connectivity, activation, operational feedback, permissible actions and assurance requirements are also applicable to the CPA.

There are cases where, instead of deploying dedicated equipment, the service provider relies on the equipment of the INP to operate service related functions; this equipment is considered shared equipment. Example cases are the configuration of classification and policing at the external interfaces of the INP border equipment for controlling the customer IP connectivity service, or the configuration of the routing and forwarding functions for constructing VPNs (see details in sections 6.1.3.2 and 6.1.6.3 respectively).

The configuration means granted to the SP are specified in the CPA. In general, configuration can be enforced using any of the following ways: a) the INP receives and implements the configuration requests from the SP in a dedicated interface, b) the equipment supports fine-grained access rights configuration, allowing the service provider to have direct access to the equipment for configuring the restricted subset of functions and physical resources, e.g. interfaces, the SP is entitled to configure, c) a configuration management system such as NETCONF [ENNS06] is employed to provide the fine-grained access rights to the shared equipment.

The SP may deploy dedicated equipment within the INP domain, as for example in the case of firewalls for providing Internet access service. In this case, the customer traffic must be routed through the firewall, with the INP enforcing the corresponding routing rule. The particular routing restrictions are specified in the CPA.

Permissible actions for controlling the connectivity provisioning applicable to CPAs are:

- shared equipment configuration
- modification of routing restrictions at the INP domain for the SP traffic
- use of network control and management functions, e.g., delegate to the INP the resource based admission control for end customer services in the case of a VoIP SP running VoIP over qualitative guarantees provided by the INP

- outsource connectivity provisioning tuning by uploading tuning rules in terms of conditions on operational feedback information to trigger tuning actions, where tuning actions may be tuning admission control parameters, or altering load balancing ratio among concurrent paths, or relaxing policing of a particular macro-flow at the expense of another macro-flow with no impact at the total traffic volume entering the INP's network, etc.

Additionally, an SP may choose to outsource the maintenance and the management of its IP equipment to the INP, making the latter responsible for firmware upgrades, performance monitoring, troubleshooting, etc.

The detailed specification of the CPA is undertaken in WP2 and will be documented in deliverables D2.1 and D2.2.

Interactions between an SP and an INP follow the same principles as the INP to INP interactions (see sections 6.1.1.2 to 6.1.1.4). The INP interconnection issues presented in section 6.1.1.5 apply as well for the connectivity provisioning offered by the INP to the SP.

6.1.3 IP Connectivity Service Interactions

6.1.3.1 Customer to Service Provider

The IP connectivity service SLA has been extensively studied in the IST-MESCAL and IST-TEQUILA projects [MESCAL, TEQUILA]. According to these studies, the IP connectivity service SLA includes among others information on the IP connectivity scope, ingress flow identification, incoming traffic conformance, performance guarantees, excess treatment and resilience guarantees (see 6.1.1.1 for a definition of these terms in the context of the INP interconnection agreement).

Within the context of an established SLA, the customer interacts with the SP to modify the capacity, to expand the scope or to upgrade the QoS and resilience guarantees of the SLA. Moreover, upon detection of a security attack, the customer should be able to notify the SP to take appropriate measures, such as modifying the SP's firewall configuration or even escalate the alarm to its affiliated SPs.

Interactions not related to the connectivity aspects of the service, such as e-mail or DNS related interactions are outside the scope of the project.

6.1.3.2 Service Provider to IP Network Provider

An SP is associated to an INP relying on the latter's IP connectivity provisioning capabilities to provide IP connectivity services to end customers. The SP buys capacity from the INP for the aggregated customer demand.

Upon the activation of an end customer service, the SP needs to configure the access termination equipment. The SP may either employ its own equipment or rely on sharing the INP's equipment to terminate the customer's access connection. In the latter case, the SP interacts with the INP to enforce the configuration at the shared equipment. Note that the cost of the access connection is covered by either the customer, e.g., to the broadband access provider it uses, the SP or the INP on behalf of the SP.

It is common for the Internet access SPs to provide protection from security threats relying on firewalls or other more sophisticated Intrusion Detection Systems (IDS). Placing a firewall in the INP's network may involve configuring the INP intra-domain routing to route the traffic from and especially towards the SP customers through the firewall. The same requirement applies in deploying IDS collectors or, in general, any complex security architecture.

In addition to offering IP connectivity provisioning capabilities to service providers, an INP may wish to exploit its IP infrastructure to offer IP connectivity services to end customers. In this case, the INP is expected to have a service provider affiliate for promoting the capabilities of its IP infrastructure, without excluding the collaboration with other non-affiliated service providers.

6.1.3.3 Service Provider to Service Provider

The interactions involved in ensuring the required restricted scope or Internet-wide IP connectivity take place between INPs. For that matter, interactions between IP connectivity SPs are not foreseen. Interactions between such SPs may occur for other purposes, e.g., in the context of a community to coordinate protection against Distributed Denial of Service (DDoS) attacks.

6.1.4 VoIP/ToIP Service Interactions

6.1.4.1 Service Provider to Service Provider

It is assumed that the interconnection between two VoIP/ToIP SPs encompasses location, signalling and media guarantees capability options.

The list of reachable telephony destinations announced by the downstream SP to the upstream SP may be fixed and specified in the SIA or it may change dynamically. In the latter case location information is exchanged dynamically. It is assumed that tier-1 SPs do exist to provide global reachability to any destination at all times. The supported capabilities, the media quality and the cost associated to each route to a telephony destination may be fixed and specified in the SIA or they may vary in time. The corresponding updates will be exchanged between the SPs through the location message exchange mechanism. A VoIP SP is expected to offer to other SPs termination to its own directly attached customers in the predefined, fixed quality and fixed cost way. For transit calls the business policy will determine the type of termination offering, fixed or dynamic.

A VoIP SP (SP0 in Figure 25) is assumed to have at least one SIA with a tier-1 SP (SP1), and multiple dynamic or fixed peering agreements with other SPs for a limited set of destinations (SP2, SP3, SP4).

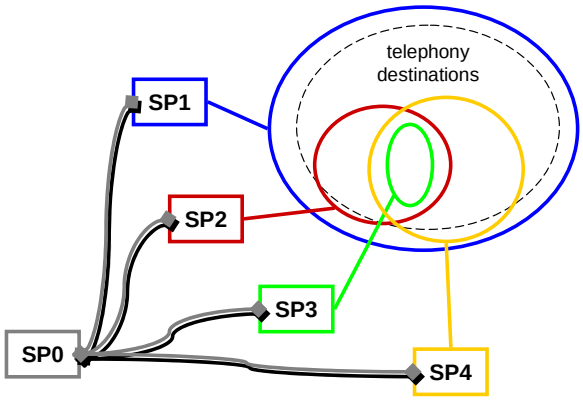

Figure 25 VoIP SPs interconnection example

In some cases it may be beneficial for an SP to route the media flow directly to the destination without relying on the downstream intermediate SPs to provide the associated media guarantees (see scenario 1 in section 4.3.2.5.1.2.2.2). It is assumed that depending on its business policies each SP may offer media guarantees as an option or mandatory, overall or depending on the telephony destination. The next hop to direct the media flow is therefore determined as either the destination IP phone, or the media gateway of the first downstream SP with a restriction for mandatory handling of the media flow.

In Figure 26 for example, SP3 implements a policy where media handling is mandatory for destinations in INP1 such as d1 and optional for destinations outside INP1 such as d2. SP2 implements a policy where media handling is optional for all destinations. As a result, the signalling for a call from SP1 will go through all intermediate providers whereas the media flow will be routed differently following the different policies of SP3 depending on the destination.

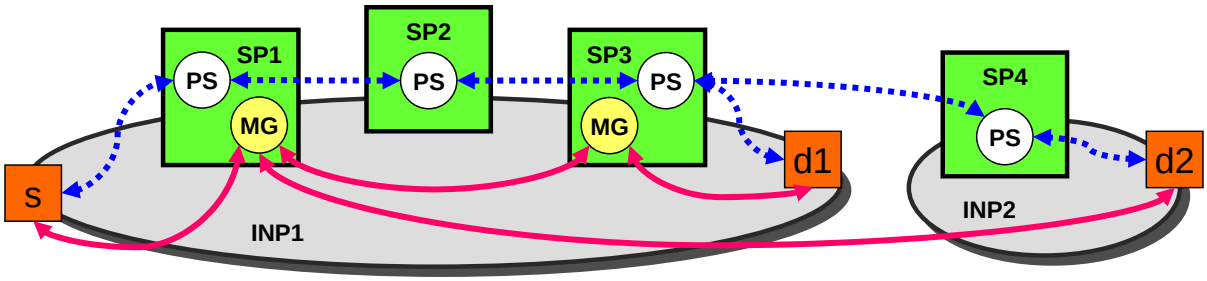

Figure 26 VoIP media flow routing example (1)

In case of optional media flow handling, the SP handling the call establishment must notify the downstream SP whether it desires to route the media flow directly to the end destination or through the downstream SP. This can be done either statically by specifying its preference in the SIA, or dynamically at call establishment time over the signalling message. Note that the next hop media gateway is not fixed, but it varies depending on the policies of the downstream SPs. As the SP needs to ensure the IP guarantees itself through its underlying INPs it must know the next hop, or a list of the candidate next hops and the associated QoS guarantees.

In Figure 27 for example, to reach the final destination SP2 may choose between SP3 and SP4. SP1 needs to know both candidates to request the associated IP QoS guarantees from INP1 towards both the SP3 media gateway in INP2 and the SP4 media gateway in INP3.

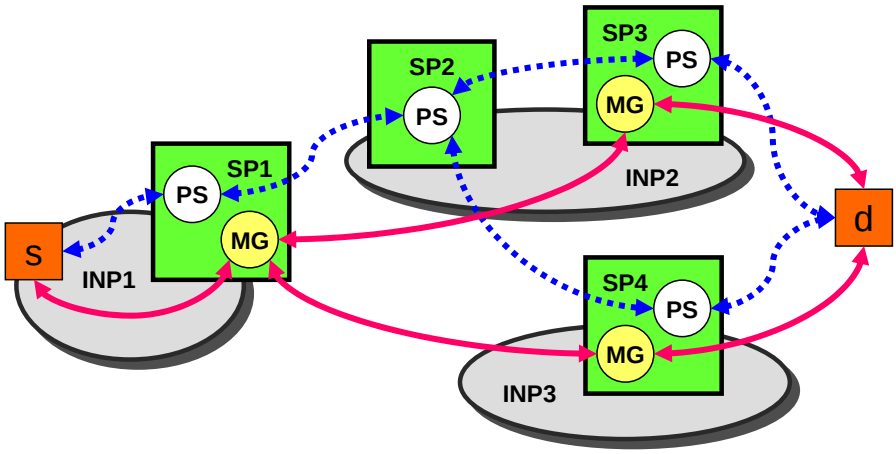

Figure 27 VoIP media flow routing example (2)

To detect INP spirals (see section 4.3.2.7.1), the call route selection process requires information on the INP path. The INP path consists of the INPs where the SP media gateways participating at the call are located (selected by the VoIP call routing selection process of the SPs), plus the intermediate INPs connecting the SP media gateway INPs (selected by the IP routing selection process of the INPs). In the example of Figure 15, the VoIP call routing selected path is {AS1, AS2, AS1, AS6}, while the IP packet routing selected path becomes {AS1, AS2, AS1, AS4, AS5, AS6}.

Information on the INPs where the SP media gateways are located is static and can be communicated as part of the route information. However, the IP paths crossing intermediate INPs change dynamically as a result of the INP routing function. Introducing a dependency between the INP routing and the SP routing processes would resolve the INP spiral problem but would considerably affect the scalability of the system. As such, interactions between the INPs and the SPs to synchronise routing should be further investigated.

Among others, an SIA will contain the following:

- statistical guarantees for the supported capabilities, quality and cost per group of fixed telephony destinations

- statistical or qualitative guarantees for the supported capabilities and quality and a cost range per group of dynamic telephony destinations
- indication of restriction for handling the media flow for fixed and potentially some dynamic telephony destinations
- activation information for the location information exchange mechanism

At bootstrapping the VoIP SPs establish SIAs. As a result the fixed telephony destinations are communicated between the interconnected SPs. In operation of the SIA, location interactions between the SPs keep the dynamic telephony destinations updated. Possible call route attributes included in the exchanged updates are:

- the supported end-to-end call signalling capabilities
- next-hop media gateway per candidate route
- IP QoS guarantees from the next-hop media gateway and beyond per candidate route (including both SP routes and direct routes if supported)
- AS numbers of the media gateways per candidate route
- tariff per candidate route

Upon each update, the call routing selection process is invoked and the optimum route is calculated based on the call route attributes and the SP's policies.

6.1.4.2 Service Provider to IP Network Provider

For the destinations that the VoIP SP aims to route the media flow directly, the SP should ensure the IP transfer guarantees itself through its underlying INPs. For the rest of the telephony destinations, the VoIP SP should only ensure IP transfer guarantees for the interconnection to the next downstream VoIP SP.

In general, the VoIP SP will issue CPA requests for either scope expansion, QoS guarantees upgrade within the existing scope, or capacity increase within the existing scope and QoS guarantees. Such CPA requests may be triggered either by the establishment of a new SIA, the reception of new candidate routes for dynamically updated destinations, the addition of a service POP, the enforcement of a new business policy or by a significant change in the call request matrix.

6.1.4.3 Customer to Service Provider

A VoIP customer is usually interested in VoIP services with universal reachability. To this end, the customer relies to a single VoIP SP to handle authorisation for signalling and financial settlement with any involved transit or terminating VoIP SPs.

Similarly to the SP to SP relationship, VoIP customers who have Internet access, i.e., an SLA with an ISP, may benefit from using it to transport the media flow of their VoIP calls, to all destinations with qualitative guarantees or to specific destinations with statistical guarantees. The translation of the specific telephony destinations to IP addresses in order to request IP QoS guarantees for, may be static or dynamic depending on the policies of the VoIP SPs.

6.1.5 Distributed VoIP Service Interactions

6.1.5.1 Service Provider to IP Network Provider

A distributed VoIP SP which relies solely on customer-operated peers (see section 4.3.2.6) is completely independent of the underlying IP network. In this case, the SP is required to employ its own mechanisms for online monitoring as it is impossible to establish CPAs with every possible INP for receiving monitoring updates for every possible peer that dynamically joins the overlay network.

When the distributed VoIP SP relies on dedicated SP-operated super-peers deployed in strategic locations with premium plane connectivity between them, it becomes dependent to the underlying IP network. In this case, monitoring updates used by the SP to exercise the overlay routing for the corresponding interconnections can be provided by the underlying INP(s).

Employing SP-operated super-peers entails establishing CPAs with the corresponding INPs, introducing administrative overhead, which is hard to scale globally with the number of INPs. Hence, unless the number of dedicated super-peers and corresponding INPs is small, the SP needs to be a large international company, able to manage agreements with multiple INPs worldwide.

Driven by the overlay network planning and provisioning functions, the distributed VoIP SP interacts with the INP to modify the capacity, to upgrade the QoS guarantees for interconnecting the super-peers or to alter the scope of the CPA to add/delete super-peers or to determine the best highly connected customer-operated peers suitable for becoming super-peers.

6.1.5.2 Customer to Service Provider

The customer of the distributed VoIP service relies on the SP for all the aspects of the VoIP service, i.e., location, signalling and media flow guarantees. Note that the SLA of the customer with the ISP only affects the first hop of the call path. Since this first next hop can be any static or dynamic peer as determined by the overlay network, customers need to have unlimited scope SLA with their ISP (see section 4.3.1). When the customer is subscribed to a quality level better than the default best-effort, there is no other implication for the distributed VoIP service besides configuring the client-side software application to use the higher quality; the mechanisms of the distributed VoIP service are capable by design to detect and exploit with fairness the varied connectivity capabilities of the customer-operated network peers.

The customer and the SP interact during the activation of the client-side software application to register the new peer into the overlay network. Then, the registered customer peer, emulating a network node, exchanges with other peers management and control information for operating the overlay network, and relays VoIP data traffic following the selected routes as determined by the logic of the overlay network engineering functions. The detailed mechanisms and information exchange at the distributed VoIP service layer need further investigation.

When the distributed VoIP service engineering uses performance monitoring to assess the connectivity of the customer peers, a dedicated function at the application needs to provide the required data. As opposed to the limited set of super-peers, it is not possible, out of scalability reasons, to acquire the required data from the underlying INP(s) for every dynamically activated peer in the Internet.

6.1.5.3 Service Provider to Service Provider

The provisioning of the distributed VoIP service itself, does not entail any inter-working with other distributed VoIP SPs. Cooperation with other VoIP SPs is required for reaching heterogeneous networks (e.g. PSTN) or closed VoIP communities (e.g. Skype). Which interconnections to establish with other VoIP/ToIP SPs is a result of the offline network planning function taking into account redundancy and geographical dispersion considerations. Online routing and traffic engineering functions operate based on the interconnections established by the offline functions.

The interactions of the distributed VoIP SP and other VoIP/ToIP SPs are similar to the interactions between VoIP/ToIP SPs (see section 6.1.4.1). The implications to the distributed VoIP functions need further study.

6.1.6 VPN Service Interactions

As part of the VPN service is to provide IP connectivity services, most of the INP to INP interactions described in section 6.1.1 also apply to SP to SP interactions provided a straightforward terminology translation: INP -> VPN SP, NIA -> SIA, ASBR -> ASBR / PE, IP address -> VPN IP address.

Regarding QoS, the VPN use case has a three-level hierarchy of actors and the QoS needs to be handled at these three layers:

- At the customer level, both end sites should be able to cooperate and agree on a coherent QoS treatment. Agreement for the interconnection between customer sites (which may belong to different VPNs in the extranet case) is out of scope of the AGAVE project but aspects of this agreement may affect the SLA of each customer to its SP.
- At the VPN service layer, all the SPs involved to interconnect two customer sites, should be able to cooperate and agree on a coherent QoS treatment. Agreements between VPN SPs are expressed in SIAs and the CPAs with the underlying INPs should reflect the SP interconnection agreements.
- At the transport layer, the INPs should be able to carry packets between VPN SP equipment at the level of QoS required by the SP.

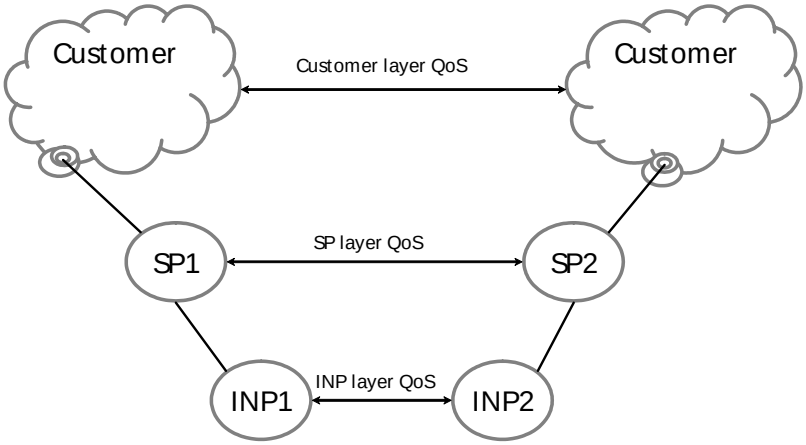

Figure 28 Three layers of QoS

Each layer should not modify the QoS identifier (IP ToS/DSCP, MPLS EXP) used by other layers, unless some DSCP remarking is explicitly requested e.g. an SP may request by the INP to remark packets between the SP and its peer SP using different DSCP values, i.e. using different DSCP at the ingress-egress identifiers.

6.1.6.1 Customer to Service Provider

A customer establishes an SLA with a single VPN SP for one or multiple VPNs (see section 4.3.3 for a description of the VPN use case). The VPN SLA may contain, among others, information on:

- VPN sites and their interconnection topology
- QoS and resilience guarantees overall or per VPN site interconnection link and for the Internet traffic
- egress policing/shaping profiles
- load balancing ratios for outbound traffic to multi-homed customer sites
- capacity sharing ratios at a VPN site interconnection link among different VPNs and/or Internet traffic
- scope and authentication information for temporary users

A non comprehensive list of the interactions (permissible actions) of the customer with the VPN SP contains the following:

- addition/removal of a VPN, addition/removal of a VPN site or site interconnection link
- increase/decrease request for QoS and resilience guarantees

- modification of the egress policing/shaping profiles; modification of the load balancing or capacity sharing ratios, user authentication configuration
- receiving fault alarms and performance analysis reports

6.1.6.2 Service Provider to Service Provider

A VPN SP relies on other VPN SPs to reach customer sites, which are not attached to its equipment. The VPN SP interacting with the customer acts as an Agent, responsible for the management of multi-party business processes, negotiations and fulfilment that allow the multi-provider VPN to function [HALS06].

Besides the customer sites directly attached to its equipment, a VPN SP may act as transit provider under the cascaded model (see Figure 17) and provide access to customer sites attached to other VPN SPs further downstream.

To support a customer-defined (partial-)mesh VPN topology, the VPN SPs need to interconnect appropriately. To this end, the VPN SP acting as Agent to the customer needs to ensure that the necessary SIAs are in place between the involved SPs. In Figure 29 for example, the customer requests VPN full-mesh interconnection between sites CE2, CE3 and C4, while CE1 is a spoke to CE2. In case SP2 acts as Agent, the existence of a suitable SIA between SP3 and SP4 must be verified. Note that the topology at the VPN SP level is independent of the customer-defined topology [NAGA04]. However, introducing additional intermediate hops is always sub optimal.

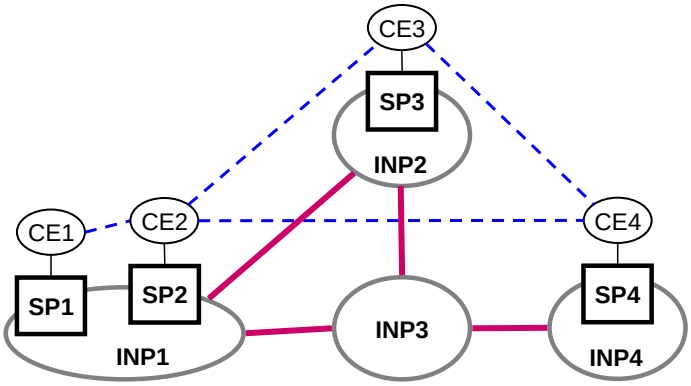

Figure 29 VPN topology example

Each SP announces its reachable customer sites (whether directly attached to its equipment or through established interconnection with other SPs), the QoS and resilience capabilities per customer site and the interconnection activation requirements, e.g. the security mechanisms or the dynamic VPN membership protocols, that need to be activated for the interconnection to be established.

The SIA for the interconnection between two VPN SPs can be used for the aggregate traffic from multiple customer VPNs. The VPN SIA may contain, among others, information on:

- union of the subsets of the customer VPN sites and their interconnection topology, applicable to the particular SP-SP interconnection
- QoS and resilience guarantees per VPN site for the aggregate traffic
- VPN specific details (egress policing/shaping profiles, etc.), relevant to the selected VPNs and VPN sites
- activation procedures at interconnection points (see interconnection scenarios in Figure 19, details in appendix section 14.2)

A non comprehensive list of the interactions between VPN SPs contains the following:

- automated VPN configuration within the context of the established SIA, e.g. VPN, VPN site or VPN site interconnection link addition/removal relying on explicit requests or on some dynamic

membership discovery mechanism inter-working between the SPs, invocation of additional capacity for a VPN site interconnection

- interactions similar to the above initiated by the VPN SP acting as the VPN agent and propagated through the adjacent VPN SPs to the VPN SPs of concern
- coordinated fault and performance management relying on the exchange of fault alarms and performance analysis reports
- authentication delegation to the agent VPN SP

6.1.6.3 Service Provider to IP Network Provider

The VPN SP relies on the INP for interconnecting its own PE and ASBR equipment as well as for interconnecting with other peer or tier-1 VPN SPs (see interconnection scenarios in Figure 19, details in appendix section 14.2).

To optimise the capacity utilisation, the paths to allocate bandwidth for customer VPNs should overlap as much as possible. This way it may be possible to benefit from the fact that the traffic leaving from a VPN site, destined to other VPN sites at anytime, is limited in total to the outbound capacity for the particular VPN customer site. To achieve this optimisation, the VPN SP should rely on *hose* scope definition in the CPA. In the hose scope there is one ingress point and multiple egress points, and the capacity is specified as the maximum guaranteed traffic volume entering the INP ingress point and leaving potentially any of the INP egress points. Optionally a maximum value (less than the one for the ingress) can be provided per egress point. This is an indication for the underlying INP to construct a tree for the egress points of the hose, which costs less than provisioning for independent point-to-point interconnection paths for each ingress-egress pair. In the example of Figure 29, INP4 can use this information to multiplex all the traffic originating at CE4 and destined to CE1, CE2 or CE3 through INP3.

The PEs and even more the ASBRs deployed by the VPN SP are in many cases equipment shared between the INP and the VPN SP. The VPN SP interacts with the INP for the management of the routing and forwarding table in the INP equipment, using the configuration means specified in the CPA.

The interconnection between two VPN SPs may span one INP domain, two neighbour INPs, or two remote INPs. The IP connectivity for this interconnection can be provided in any of the following ways: a) one SP requests and administers the bi-directional interconnection between the two SPs by the INP hosting its equipment, b) each one of the two interconnected SPs requests the uni-directional connectivity from the INP for the traffic leaving its equipment towards the peer SP equipment, c) both SPs request and administer a multi-party CPA, where both SPs jointly administer the CPA offered by the cooperation of the underlying edge INPs, e.g. for interconnecting SP2 to SP4 in Figure 29, INP1 and INP4 would be involved. The latter option is not taken into consideration in the current CPA specifications.

6.2 Network Plane Engineering

6.2.1 Motivation

In order to support the demanded QoS requirements from Service Providers (SPs) in a cost efficient way, IP Network Providers (INPs) need to properly allocate and engineer their network resources. DiffServ has been regarded as a promising control-plane paradigm for achieving service differentiation through differentiated forwarding mechanisms (Per Hop Behaviours (PHB) and Per Domain Behaviours (PDB)). But there are other complementary means of supporting service differentiation. A typical example is to route traffic with different QoS requirements through distinct paths that are able to satisfy the specific performance demands of the traffic. From the viewpoint of INPs, this type of routing differentiation not only supports heterogeneous QoS requirements, but it is also a useful tool for resource optimisation purposes such as load balancing, as traffic flows may

follow different paths to reach the same destination with/out coherent/similar IP treatment, or load sharing. We use the term *Network Planes* to refer to the distinct resource provisioning paradigms that apply the aforementioned mechanisms for supporting intra-domain QoS differentiation and Traffic Engineering (TE). These Network Planes are *internally* implemented by the INP so as to convey IP traffic related to QoS-aware services offered by SPs. The interconnection of adjacent NPs can lead to the emergence of Parallel Internets (PI) suitable for dedicated IP-based services.

6.2.2 Network Plane definition

A Network Plane is a logical partition of an INP's network resources designed to transport traffic flows from services with common connectivity requirements. The traffic delivered within a Network Plane has particular treatment in both forwarding and routing so that service differentiation across NPs is enabled in terms of edge-to-edge QoS, availability and also resilience.

NPs can be used to convey traffic from one or several services managed by one or more SPs. They may be proactively engineered to meet anticipated service requirements of end customers/SPs or they may be built/reconfigured to meet actual service requirements raised by a SP. INPs and SPs agree on how traffic associated with a particular CPA will be identified but the mapping to Network Planes is an internal matter for an INP, and not visible to SPs.

The realisation of NPs may be distinguished according to the following aspects:

- *Routing aspect*: IP packet treatment is differentiated in terms of the routes taken through the INP's network. The options for differentiated routing include:
  - *Dedicated topology*: A NP is formed by a dedicated physical or logical topology.
  - *Dedicated routing selection process*: NPs may be distinguished according to the route selection processes they adopt. Multiple route selection processes may operate on the same topology or a different route selection process can be allocated for each topology.
  - *Different routing convergence times*: The differentiated routing can be result of the means that are deployed to enhance the stability of the routing tables;
  - *Different fast reroute procedures*: When errors or failures occur in a given topology, the routing process can be enhanced by means to fast-reroute the IP traffic;
  - *Different policies and metrics*: Another alternative to implement differentiated routing is to have dedicated metrics, such as MT-OSPF link weights, for each NP.
- *Forwarding aspect*: At the forwarding level, an INP can deploy different queuing and scheduling mechanisms, e.g. DiffServ PHBs, for traffic belonging to different NPs.
- *Resource Management aspect*: The treatment experienced by IP packets can be differentiated by the admission control, traffic shaping and policing mechanisms associated with a NP and the resources allocated to it in terms of dedicated or shared network capacity.

6.2.3 Network Plane creation and realisation

As mechanisms such as DiffServ and multi-topology routing protocols are not deployed everywhere in the Internet, Network Plane realisation should be flexible enough to accommodate various scenarios. Two examples of the NP realization means are described in the following subsections: one purely based on differentiated forwarding, the other on differentiated routing. It is expected that in the general case INPs with adopt the more sophisticated strategy of realising NPs based on more than one mechanism simultaneously.

6.2.3.1 Network Plane realisation through differentiated forwarding

The DiffServ model was designed with the philosophy of pushing complex processes and resource management to the edge of networks while keeping the packet handling in the core simple. Instead of focusing the QoS problem on a per flow basis as previously done by the IntServ model, a reduced set

of QoS classes was proposed as a way to achieve scalability. There is no more connection state in the core network, and forwarding behaviour would be based on packet marks indicating the class of service. This forwarding behaviour is called Per Hop Behaviour (PHB), and consists of treating packets in each hop according to the QoS class they belong to. The differentiated treatment of packets is fundamentally based on the usage of scheduling algorithms that serve different classes with different priorities. The DiffServ paradigm assumes that traffic injected in the network is controlled in the edges of it by means of traffic policing/shaping mechanisms.

In practice, network operators who desire to offer some kind of QoS guarantees must establish the appropriate network parameters (for the PHB mechanisms and for the traffic policing/shaping mechanisms) and allocate the necessary network resources to achieve the desired Per Domain Behaviour (PDB). In the AGAVE project, this is called "over-provisioning factor" associated with individual network planes to achieve dedicated QoS requirements.

6.2.3.2 Network Plane realisation through differentiated routing

In today's Internet, including those domains that have already deployed DiffServ implementations, packets belonging to the same traffic class always follow the same physical path (or paths in the case of ECMP) towards the destination. This is because the routing protocols used by most INPs, such as conventional OSPF and BGP, use a single Routing Information Base (RIB) for all types of traffic. However, differentiated routing can be also used for the realisation of Network Planes. Typically this can be achieved through (multi-topology enabled) IP routing, MPLS explicit routing, overlay routing, etc. In this scenario, traffic belonging to different Network Planes may use different routes towards the same destination.

In order to achieve this, multiple strategies can be envisioned:

- One routing protocol dealing with multiple NPs. In this way, a dedicated RIB may be maintained for each Network Plane. In fact, most IP based routing protocols such as OSPF, IS-IS or BGP have already been extended to be multi-topology aware, and their extensions are known as Multi-topology OSPF/ISIS/BGP (M-ISIS [PRZY05], MT-OSPF [PSEN05]) and Multi-protocol BGP (MP-BGP) [BATE02] respectively. Currently, these protocol extensions are often used for carrying *different types* of traffic in different planes, e.g., unicast/multicast or IPv4/IPv6 traffic. However, this functionality is also expected to be used for other purposes such as QoS service differentiation across multiple Network Planes. Another option is to keep an only one RIB that keeps information on different routes for reaching the same destinations, so that in function of the Network Plane and its current state (in terms of load), a given route would be chosen;
- Multiple instances of the same routing protocol for different NPs. Instead of adapting the protocol to deal with multiple topologies, one instance of a simple protocol (such as OSPF or IS-IS) may be run for each NP. Logically, if the protocol is the same, to achieve routing differentiation it is necessary to have different topologies on each Network Plane. Implementing this option with real Network Elements would consist on selecting which instance has to solve the routing, packet by packet after determining which NP the packet belongs to.
- Different routing mechanisms for different Network Planes. This in an obvious extension of the previous case, where instead of running the same routing protocol at each Network Plane, different ones are run at every Network Plane. In this way, it is not strictly necessary that the topology of each Network Plane be different.

It should be notes that the above routing strategies do not only apply to IP based routing protocols, but also to other routing mechanisms such as MPLS based explicit routing and various overlay routing approaches.

A more detailed discussion on routing issues for supporting Network Plane engineering is covered in section 13

6.2.3.3 Network Plane Engineering Design Challenges

6.2.3.3.1 Stability

In today's best effort based Internet, the stability issue often refers to BGP-based routing and also IGP protocols in case of network failures (e.g., OSPF convergence). As far as Network Plane engineering is concerned, additional stability issues need to be considered, especially when multiple Network Planes are engineered in an integrated fashion. Typically, as traffic is allowed to be switched to alternative Network Planes so as to maintain QoS performance or avoid congestion, chances are that new congested points are formed at some physical links due to *uncoordinated* Network Plane switching. As a result, traffic may turn to other Network Planes (even back to their original NPs) after detecting the new hot spot. Hence, this type of large-scale path switching/oscillation may cause serious stability problems.

6.2.3.3.2 Scalability

Scalability is another vital factor to be addressed in Network Plane engineering. One crucial issue to be investigated by each INP is - what is the most appropriate number of Network Planes to achieve the best trade-off between QoS differentiation and scalability? For instance, if multi-topology routing protocols are used, the routing infrastructure needs to be scalable in (1) the size of the overall routing table for all NPs and (2) the number of control packets such as MP-BGP updates to be advertised. If MPLS is used for implementing Network Planes, the number of LSPs to be established is also a key issue in scalability. Apart from the routing aspect, scalability issues also exists in forwarding mechanisms such as DiffServ, where the total number of QoS Classes (PHBs) and DSCP allocation are the major consideration by INPs that apply the DiffServ paradigm.

6.2.4 Using Network Planes

As it has been mentioned, Network Planes are engineered for service differentiation purposes, for load balancing and for engineering robust and resilient networks. In this section we illustrate how Network Planes are manipulated for such purposes. The objective is that, the underlying network is highly utilised while packets in individual Network Planes are effectively treated according to the QoS requirements specified by the services.

6.2.4.1 NPs for service differentiation and QoS provisioning

In principle, DiffServ is deployed where the Internet bottleneck is located, i.e., prioritised forwarding is strictly necessary where bandwidth resources are not sufficient for delivering all the traffic without congestion. This may happen no matter whether the bottleneck is at the core network or the edge (e.g., inter-domain links) or both. From an objective perspective, QoS service differentiation in the DiffServ environment normally includes such metrics as (queuing) delay, packet loss rate and jitter. On the other hand, many INPs currently tend to adopt the strategy of over-provisioning their network resources, typically at the core of the network. In this case, the standard DiffServ paradigm is not needed. However, this does not mean that service differentiation is no longer applicable. Indeed, even in an over-provisioned network, service differentiation can be still useful when Network Planes are engineered with other objectives in mind, such as load balancing. Figure 30 illustrates a simple example of end-to-end delay differentiation in an over-provisioned domain running multi-topology IGP. The purpose of this configuration is load balancing across the physical network. It is assumed that three types of *delay-sensitive* services from S to D are mapped onto three network planes, each of which has a dedicated set of IGP routing metrics. As the network bandwidth is over-provisioned, the requirement of zero packet loss from the three services can be automatically satisfied. On the other hand, service differentiation in end-to-end delay can be achieved through delivering the traffic via paths with different hop counts. This is because propagation delay becomes the sole factor that influences the end-to-end delay performance, as queuing delay is also negligible in an over-provisioned network. As shown in Figure 30, three traffic classes are delivered through three non-overlapping routes decided by dedicated MT-IGP link weights for each NP. The advantage is that, the

network can be more load-balanced without violating the different delay requirements from individual QoS classes. Otherwise, if all traffic uses the same path with minimum end-to-end delay, e.g., the direct link from S to D, then this link has higher possibility to be over-loaded and the network becomes less balanced.

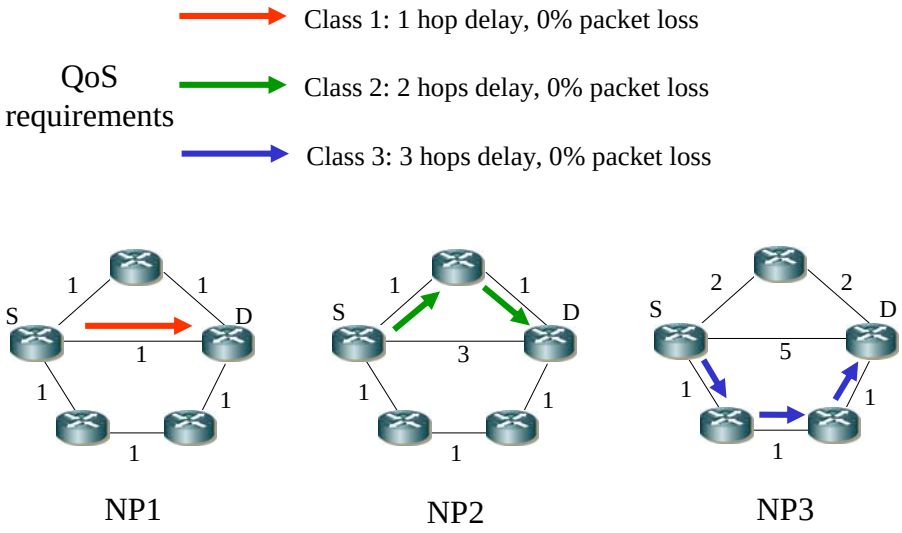

Figure 30 Service differentiation through routing

6.2.4.2 NPs for resource optimisation and load balancing

As we have previously mentioned, Traffic Engineering, typically load balancing and load sharing, are crucial issues to be investigated when engineering Network Planes. In today's IP based network, load balancing is often achieved through multi-path routing and load sharing mechanisms such as Equal Cost Multi Path (ECMP). Recently, offline IP based TE algorithms have been proposed to calculate optimal IGP link weights and BGP route attributes for load balancing purposes both within one domain and across multiple ASes [FORT02] [QUOI03]. Apart from the conventional IGP/BGP based solutions, MPLS explicit routing and overlay routing can also be used for Traffic Engineering purposes. When the physical network is engineered into NPs, the above TE approaches can be also applied. In fact we can envisage two different strategies of network plane TE: (1) per plane based TE and (2) Integrated TE across NPs.

6.2.4.2.1 Load Balancing within each Network Plane

One straightforward strategy to achieve load balancing is to directly extend the aforementioned solutions to multiple Network Planes individually. Taking intra-domain routing as an example, ECMP can be enabled within each Network Plane and dedicated IGP link weights can be optimised based on the forecasted traffic demand for each NP. The advantage of this approach is that existing IP based Traffic Engineering solutions can be directly borrowed and used. However, since traffic from all Network Planes is mapped onto the same physical network infrastructure [1], *independent* TE paradigms within each NP might not be efficient even if efforts are made for traffic optimisation within individual planes.

6.2.4.2.2 Load Balancing across Network Planes

For this integrated strategy, Traffic Engineering is performed by taking into account the traffic from all the Network Planes at the same time. In this case the INP is able to coordinate NPs for better traffic

[1] In some cases, this might not be the case as a network link may be dedicated to only a subset of network planes, e.g., for the premium traffic plane. In this scenario, the view of the logical network topology for each Network Plane is different.

distribution. This type of TE interaction across multiple NPs within a domain will be described as *vertical* Traffic Engineering in the AGAVE project. In realising vertical TE, more sophisticated optimisation solutions are needed compared to per NP Traffic Engineering.

Apart from service differentiation, Network Planes can be also created for Traffic Engineering purposes. In this scenario, multiple *equivalent* Network Planes are generated for each single class of service. As a result, traffic with the same QoS requirements can be "transparently" delivered through multiple equivalent NPs with different path selections. In time of network congestion, traffic can dynamically switch onto alternative equivalent Network Planes to avoid traffic concentration. This can be regarded as "vertical" load sharing across multiple Network Planes.

6.2.4.3 Robustness and resilience in NPs

In general, there are two major issues to be considered in the robustness issue of NP engineering. The first issue is the resilience of the network/service to failures. In case of link/node breakdown, the traffic needs to be delivered through alternative routes with minimum impact on the current QoS and TE performances. The second issue to be considered is the resilience to external traffic dynamics. Specifically, the traffic pattern for a network varies from time to time, which means that a single static network configuration might not be able to effectively handle the dynamic traffic demand all the time. In general, there are two complementary solutions to this problem. The first solution is to provide efficient online or reactive TE mechanisms for coping with external traffic dynamics. That is, the network configuration is dynamically (ideally in a local scope) changed according to the monitored change of traffic pattern. The second solution is to come up with a single offline network configuration, by means of pro-active TE, that performs well for multiple distinct traffic patterns that can be forecasted. This network configuration might not be *optimal* for any specific traffic pattern, but it is able to produce *satisfactory* performance for all the predictable traffic matrices.

As far as Network Plane engineering is concerned, both the above-mentioned approaches can be applied for multiple NPs. Similar to the description in section 6.2.4.2, robustness issues can be considered either on per NP basis or simultaneously across Network Planes. In the first case, conventional approaches are "replicated" to all Network Planes so that each of them applies *independent* resilience-aware configuration. In the second scenario, robustness is considered in an integrated fashion across multiple NPs. For example, traffic can be automatically delivered through alternative paths *provided by other "compatible" Network Planes* if the original NP has encountered network failures or congestion.

The main metrics to measure and reflect the level of resilience of IP network or even IP services in general, are the MTTR (Mean Time To Repair), MTBF (Mean Time Before Failure) and the availability, see also Section 12. Agreements like SLA, NIA, SIA, may include or exclude, in addition to the aforementioned metrics, the maintenance windows defined by the SP. The following section provides some useful definitions so as to understand the notion of resilience in IP network. The purpose of this section is not to provide in depth, information about resilience and means to measure it, but only to provide some useful definition. This terminology will be exploited during the specification of the algorithms and protocols during the WP3 design phase.

6.3 Performance and Traffic Monitoring

6.3.1 State of the art

Numerous Working Groups (WGs) have been chartered within the IETF to research on monitoring and measurements issues. These WGs have defined (and *still are defining for some of them*) metrics such as OWD (One Way Delay) or OWDV (One Way Delay Variation), promoted common IP traffic flow measurement scheme and specified a set of capabilities for sampling packets. Below a list of IETF WGs focusing on monitoring and measurement issues:

- In 2000, the IETF Real Time Flow Measurement Working Group (RTFM) concluded after delivering several RFCs defining a mechanism to configure network traffic flow meter devices

and collecting the flow data from remote meters. Since then, several IETF Working Groups have worked on network performance monitoring.

- IP Performance Monitoring WG (IPPM) has issued several RFCs to define different metrics: metrics for measuring connectivity [MAHD99], one-way delay metric [ALME99a], one-way packet loss metric [ALME99b], round trip delay metric [ALME99c], one-way loss pattern sample metric [KOOD02] and one-way IP packet delay variation metric [DEMI02], as well as means to measure these metrics with periodic streams [RAIZ02].
- IP Flow Information Export WG (IPFIX) has worked on the definition of a standard way of exporting information that may be relevant for accounting, Traffic Engineering, QoS monitoring and any other applications related to IP flows. In the context of IPFIX, an IP flow is defined as a set of IP packets passing an observation point in the network during a certain time interval and that have a set of common properties, ranging from the next hop IP address (i.e., all packets outgoing an interface) to the transport or application header field (destination port number or for instance RTP header field). Derived from Cisco NetFlow protocol, the resulting IPFIX protocol defines how IP flow information can be exported from routers, measurement probes or other any other devices. However it does not define the information format neither the information to exchange, so that it might be used in different contexts.
- Packet Sampling WG (PSAMP) has defined packet sampling techniques required for measurement applications. The packet sampling protocol provides mainly mechanisms for the packet selection using different filtering and sampling techniques, as well as a standard way for the export and storage of the sampled packet data. PSAMP relies on IPFIX architecture for the export of packet records that can be seen as a special IPFIX flow record.
- Remote Network Monitoring MIB WG (RMONMIB) has edited a large set of RFCs that aim to define the interface between a monitoring agent (i.e. a probe or more probably a network element software component) and a management application. To this purpose, a number of SNMP managed objects have been defined to provide the ability to manage fault, configuration and performance in multiple network layers. Among the numerous MIBs defined by RMONMIB, the following ones can be listed: RMON MIB [WALD00] defines objects for layer 2 monitoring and RMON-2 MIB [WALD99] enables to aggregate monitoring per transport or service layer ; DSMON MIB [BIER02] enables to aggregate monitoring per DSCP value; APM MIB [WALD04] enables to aggregate monitoring per application. Finally, the RAQMON MIB enables to report end-to-end QoS experience for real time applications such as VoIP.
- Prefix Taxonomy Ongoing Measurement and Inter Network Experiment WG (PTOMAINE) was chartered to consider and measure the problem of routing table growth and possible interim methods for reducing the impact of routing table resource consumption within a network and the global Internet. This WG did not focus on data itself but on the measurement of routing tables. Only one RFC was edited by this WG, RFC3765 [HUST04].

Beside this standardisation effort, collaborative projects have been launched to deal with monitoring and measurement issues. Hereafter a list of IST projects that studied the aforementioned issues:

- IST-INTERMON project [INTERMON] aimed to enhance the inter-domain QoS and traffic analysis in large-scale, multi-domain Internet infrastructures. INTERMON architecture was based on a distributed information base which collects data provided by both IPFIX measurement tools and analytical models. The main output from INTERMON project was the implementation of visual data mining techniques to ease the analysis of the information collected about network such as network topology, traffic and Quality of Service and to verify the SLA fulfilment in both intra-domain and inter-domain scenarios.
- IST-LOBSTER [LOBSTER] aimed to deploy an advanced pilot European Internet Traffic Monitoring Infrastructure based on passive monitoring sensors at speeds starting from 2.5Gbps and possibly up to 10Gbps. LOBSTER ambition was to enable full cooperation among points of presence, and contribute towards effectively monitoring the underlying network, providing early

warning for security incidents, and providing accurate and meaningful measurements of performance.

- The IST-6QM [6QM] focuses on IPv6-related measurement technologies for Quality of Service. The project main goal is to develop a system with the required functions for QoS measurement, such as packet capturing, precise time-stamping, data collection, QoS metrics derivation and result presentation.
- IST-SCAMPI [SCAMPI] project goal was to develop a scaleable monitoring platform for the Internet. It also aimed to promote the use of monitoring tools for improving services and technology.
- IST-MOME [MOME] main objective was to co-ordinate activities in the field of IP monitoring and measurement by offering a platform for knowledge, tool and data exchange.

6.3.2 Inter-domain monitoring and measurements challenges

Monitoring and measurements architectures are currently deployed within domains of individual providers. The scope of the monitored information is limited by the boundaries of a single provider/domain. Performance and traffic-related data are judged critical information. Some INPs are publishing in real time this data, like AT&T at the web site http://www.att.com/ipnetwork. But other INPs are still reluctant to publish this sensitive information mainly due to strategic concerns such as competing INPs might exploit performance related information of competitors to strengthen their own position in the market.

Beside these strategic considerations, the cooperation of several INP to provide global services will promote and accelerate the need of sharing monitoring and measurement data so as to evaluate the level of the Quality of Service experienced by flows associated to QoS-enabled services. Thus, technical issues should be handled so that to deploy an inter-domain/inter-provider monitoring and measurement architecture as a companion means for the deployment of inter-domain services. Hereafter a list of constraints that should be considered:

- *Monitoring scope*: When several providers are collaborating in order to provide monitoring information, the scope of the monitored network must be specified. Particularly, adjacent providers must decide if the inter-domain link is part of the scope of monitored network, and determine who will take in charge the inter-domain link and how it will be monitored;
- *Metrics and Measurement methods*: Providers must agree on the metrics that will be monitored, the strategy of measurement, the interval of measurement, how to conduct measurement, data format, types of alarm, how alarms are generated, etc.
- *Synchronisation*: Measurement information should be synchronised so measurement output can be valid and meaningful for all collaborating providers;
- *Exchange of measurement data*: Providers should agree how to exchange the collected data and how it will be aggregated/concatenated so as to have an e2e meaning.
- *Suitable for both centralised and cascaded peering model*: The monitoring and measurement architecture should be suitable for both centralised and cascaded models.

Within the AGAVE project, the collected measurement data is used for the following purposes:

- *Agreement assurance*: Monitoring functions are mandatory to assess the terms of the different kinds of agreements that are defined in AGAVE context: SLA, SIA, CPA and NIA. Two levels of monitoring have to be distinguished: service monitoring and network monitoring.
    - *Service monitoring*: Service monitoring undertaken by service providers typically aims to provide some real time information about the status of the service they provide. The parameters that have to be monitored have to correspond to the characteristics that are specified within the SIA agreements that they have with the other Service Providers they are peering with and to the characteristics that are specified within the SLA agreements

that they have with their customers. These characteristics depend on the type of services that are offered by the Service Providers and for each type of service, different elements might be monitored, such as signalling (i.e., call management for VoIP services) and different kinds of data (i.e. digital voice transport for VoIP services).

- *Network monitoring*: Network monitoring typically aims to provide INPs with some real time and non real time information about the status of the Network Plane they are managing and of the traffic classes they provide to the service providers. The parameters that have to be monitored for traffic classes have to correspond to the characteristics that are specified within the CPA agreements that they have with the Service Providers to which internet network providers offer these traffic classes. The parameters that have to be monitored within a Network Plane have to correspond to the characteristics that are specified within the NIA agreements that they have with the other Internet Network Providers that they are peering with. These characteristics are different for each Network Plane and depend on the way the Internet Network Provider assigns the Service Providers traffic classes onto each Network Plane. However, within a given network plane, no distinction can be made between traffic classes. Network Plane monitoring therefore applies for any traffic transported by a given network plane.

- Traffic Engineering: Traffic engineering [AWDO02] is a network engineering process that aims at evaluating and enhancing operational IP networks performance. INPs would typically use Traffic eEngineering to meet the CPAs and NIAs agreed with SPs and peer INPs. Similarly SPs would use traffic engineering to meet the SLAs and SIAs agreed with their customers and peer SPs. In fact, Traffic Engineering related monitoring may be provided as part of the connectivity provisioning offering of the INP to the SP, thus enabling the SP to exercise Traffic Engineering on the IP resources it runs its services over. Traffic Engineering is often opposed to over provisioning, i.e., providing enough networking resources (typically link capacity) for sustaining any predicted growth in network resources. Typical examples of Traffic Engineering objectives are avoiding congestion, spreading the traffic across the network so as to balance the load and minimizing the vulnerability to outages (errors and failures). The Traffic Engineering process often relies on a closed-loop optimisation process which takes the current operational state of the network as input. This state can only be determined by measuring the network. The metrics as well as the measurement granularity depend on the traffic engineering objectives. Different Traffic Engineering objectives are applied to different Network Planes, depending on the traffic classes assigned to each network plane. The operational state of each Network Plane can be measured independently. However, the different Network Planes cannot be engineered in an independent manner since the IP network resources are shared between the Network Planes.

Note that these measurement and monitoring data can be used for performance analysis since business and network planning and dimensioning processes can greatly benefit from information on how well the network fits the actual demand, how well control and management processes behave to meet the target policies, how the demand from customers and peers is actually formulated against predictions derived from theoretical market analysis, etc. Measurement data can be used to derive this kind of information to enable the conduction of what-if scenarios, to provide tools to facilitate the long-term business and network planning processes.

6.3.3 Measurement Metrics

Metrics that have to be monitored depend on the kind of relationship entities have with each other: INP providing connectivity to an SP or another INP, SP or INP sending traffic towards an INP, SP peering with another SP, SP providing services to customers.

6.3.3.1 IP metrics

The following is a non exhaustive list of metrics required for the provisioning of QoS and resilience over an IP network:

- *Effective load/throughput per SLA/SIA/CPA/NIA*: This metric can be used to evaluate the usage of the corresponding agreement, either for agreement assurance, traffic engineering or performance analysis purposes;
- *Effective Load/throughput per network plane*: This metric is exploited by the INP for Traffic Engineering and performance analysis purposes;
- *End-to-end load associated with inter-domain QoS-inferred routes*: this metric is used for Traffic Engineering purposes and can drive the inter-domain route selection process;
- *Actual load of LSP/tunnel*: this metric can be used for traffic engineering purposes;
- *Actual load of edge/core and ASBRs routers*: this metric can be used for Traffic Engineering purposes in order to avoid congestion nodes;
- *OWD (One-Way packet Delay)* of intra-domain/inter-domain routes/AS hops;
- *OWDV (One-Way packet Delay Variation)* of intra-domain/inter-domain routes/AS hops;
- *OWL (One-Way packet Loss)* of intra-domain/inter-domain routes/AS hops;

For interpretation purposes of the aforementioned metric, the "Diameter of the AS/domain" should also be taken into account.

6.3.3.2 Indicative IP performance objectives

ITU-T Recommendation [Y.1541] suggests network performance objectives for different kinds of QoS traffic characteristics. Class 0 is intended to real-time jitter-sensitive and highly interactive applications such as VoIP, Class 1 is intended to real-time jitter-sensitive application, Class 2 is intended to highly interactive transaction data such as signalling messages, Class 3 is intended to transaction data applications, Class 4 is intended to low loss applications such as video streaming, and Class 5 is intended to default IP applications.

These objectives are expressed in terms of network performance parameters defined in ITU-T Recommendation [Y.1540]: IPTD (IP Packet Transfer Delay) that corresponds to OWD (One-Way packet delay, IPDV (IP Packet Delay Variation) that approaches OWDV (One-Way packet Delay Variation), IPLR (IP Packet Loss Ratio), that corresponds to OWL (One-Way packet Loss) and IPER (IP Packet Error Ratio), that additionally takes into account error received packets. All these parameters are supposed to be computed during a given measurement period with a given tolerance interval and only when IPSA (IP Service Availability function) attests that IPLR is above a given threshold.

The following table is derived from [Y.1540] ITU recommendation and is indicative of the IP traffic performance metrics and objectives to be used in AGAVE.

Network performance parameter	Nature of network performance objective	QoS Classes					
		Class 0	Class 1	Class 2	Class 3	Class 4	Class 5 Unspecified
IPTD	Upper bound on the mean IPTD (Note 1)	100 ms	400 ms	100 ms	400 ms	1 s	U
IPDV	Upper bound on the $1 - 10^{-3}$ quantile of IPTD minus the minimum IPTD (Note 2)	50 ms (Note 3)	50 ms (Note 3)	U	U	U	U
IPLR	Upper bound on the packet loss probability	1×10^{-3} (Note 4)	1×10^{-3} (Note 4)	1×10^{-3}	1×10^{-3}	1×10^{-3}	U
IPER	Upper bound	1×10^{-4} (Note 5)					U

Table 1 Y.1541 – Provisional IP network QoS class definitions and network performance objectives

6.4 Data plane functions

In order for INPs to deploy and operate Network Planes a set of features are required in the data plane to identify the traffic belonging to different NPs and to apply appropriate forwarding behaviours to the packets.

The data plane functions of concern are: *Packet Forwarding*, which encompasses egress interface selection and queuing/scheduling mechanisms; and *Traffic Conditioning* which consists of packet classification, metering, marking and shaping/dropping.

One example of the use of these functions is specified by DiffServ. Packets are marked with a DiffServ codepoint (DSCP), which is stored in the ToS octet in IPv4 or the Traffic Class octet in IPv6. Although AGAVE is not exclusively dependent on the PHB queuing/scheduling techniques of DiffServ for implementing NPs and enforcing QoS, all of the DiffServ data plane functions are assumed to be available for INPs to engineer NPs.

The exception to the above, and possibly the most important function for AGAVE, is the egress interface selection aspect of packet forwarding. As discussed in section 6.2.2, there are several options for implementing NPs based on differentiated routing mechanisms: multi-topology routing, multiple instances of routing protocols or utilisation of different routing protocols per NP. None of the multi-topology routing proposals to date prescribe the means by which traffic should be identified as belonging to a particular topology. In AGAVE we assume a similar approach to DiffServ where the ToS/Traffic Class octet is used to identify to which NP a packet belongs while it is routed within the boundaries of an INP. It is assumed that routers will have the functionality to index the appropriate FIB table/entry based on the contents of the ToS/Traffic Class octet and therefore forward packets to different egress interfaces depending on the NP they belong to and the prior routing decisions made by the routing protocol operating in that NP.

As in DiffServ, packets may be marked/remarked at the ingress ASBR of an INP's domain and also at the egress ASBR to downstream INPs. The former case is to exchange the external packet marking agreed in the NIA/CPA with upstream INPs/SPs with the internal marking used for internal NP identification. The latter case is to map internal NP markings with those agreed in NIAs with downstream INPs. Traffic shaping may be optionally deployed at both ingress and egress of an INP's domain. It should be noted that packet marking on inter-domain links is specified within the NIA and has only local significance. It does not imply any particular NP implementation technology within the downstream INP, e.g., DiffServ.

If an INP is using both DiffServ PHBs and differentiated routing to engineer NPs then there is a potential clash with the use of the ToS/Traffic Class octet. It is an internal matter to an INP to determine whether the ToS/Traffic Class octet is partitioned into DCSP and MT-ID fields so that the routing plane and PHBs experienced by packets is controlled independently or whether a NP enforces both routing and queuing/scheduling with a common codepoint.

As far as inter-domain is concerned, further consideration is needed. In MP-BGP [BATES02], each routing topology is identified in the combination of the address family identifier (AFI) and sub-AFI(SAFI). In this case, the reachability information across multiple domains is advertised on per AFI/SAFI basis. In order to enable inter-domain routing differentiation, a proper mechanism is needed to map the DSCP/MT-ID values to AFI/SAFI for supporting Parallel Internets routing. If q-BGP is used for QoS-aware inter-domain routing, the QoS Class (QC) identifier should also be handled in a similar fashion.

As an alternative to using IP-layer packet forwarding mechanisms an INP may decide to use MPLS LSPs to engineer its NPs. In this case the data plane functions involved are the classification aspects of traffic conditioning, in terms of determining the forwarding equivalence class of ingress traffic and hence the LSP to which it should be forwarded, and the standard label switching and stacking functions of LSRs/LERs. At this stage, no additional data plane functions beyond those currently deployed for MPLS are foreseen by AGAVE.

6.5 IPv6 and Multicast

IPv6 and Multicast are outside the scope of the AGAVE project.

6.6 Building Parallel Internets

The concept of Parallel Internets is introduced by AGAVE as an innovative way to enable *end-to-end* service differentiation across multiple administrative domains. Specifically, Parallel Internets are coexisting parallel networks composed of interconnected per-domain Network Planes. Parallel Internets are constructed from the perspectives of each INP, by configuring for each Network Plane different inter-domain routes to certain destinations, based on local criteria. For each Network Plane, traffic may enter/exit the INP domain through a different AS Border Router (ASBR), or through different portions of the same inter-domain link (e.g. based on DiffServ capabilities). How to horizontally build Parallel Internets with individual Network Planes is through the negotiation by individual autonomous domains, typically via the establishment of NIAs.

In a similar way to the engineering of Network Planes, NP binding for constructing PIs can be also performed with multiple dimensions, namely by means of routing and/or forwarding differentiation and resource provisioning. Compatible Network Planes configured by individual domains (e.g. with similar QoS characteristics and requirements) are horizontally bound together to extend the corresponding QoS services to each other. In the data plane, this can be achieved using dedicated DiffServ code points (DSCPs) agreed in the established NIAs negotiated by neighbouring domains. In the dimension of routing differentiation, independent inter-domain routing can be also specific to individual PI planes. This can be achieved by introducing new means ensuring coherence and consistency of treatment when crossing several INP domains such as MP-BGP [BATE02], q-BGP and meta-QoS classes. In this case, one or multiple inter-domain routing topologies can be dedicated to serve each PI plane according to the required QoS performances. If MP-BGP is used for constructing Parallel Internets, the routing topology for each plane is identified by the combination of the Address Family Identifier (AFI) and Sub Address Family Identifier (SAFI). Similarly, if q-BGP is used, inter-domain routing within each plane of Parallel Internets can be identified by the dedicated QoS-class identifier. If meta-QoS classes are used, each INP classifies its Network Planes according to the meta-QoS classes desired (i.e. set of classes used to classify Network Planes designed to support services with similar constraints, resilience and traffic protection means) so that NPs from different INPs, but belonging to the same meta-QoS class, can be bound. Multiple dimensions of the abovementioned approaches can be applied in combination, and this is usually negotiated by neighbouring domains with NIAs.

In case that some administrative domains do not support Network Planes, Parallel Internets can be only built on top of the local best effort service. In this case, service differentiation in this particular domain is not possible in either the forwarding differentiation or routing differentiation. As a result, from an end-to-end perspective, quantitative service differentiation or QoS guarantees may not be feasible in general. The only possible scenario for a successful quantitative service differentiation is that the non-NP-aware domain is over-provisioned such that it never becomes the bottleneck of the Parallel Internets. Nevertheless, *qualitative* service differentiation may still be possible even if an under-provisioned and non-NP-aware domain is involved. Provided that the bottleneck domain of the Parallel Internets maintains and properly engineers Network Planes, it is always possible to achieve some certain level of end-to-end service differentiation across multiple domains.

On the other hand, IP tunnelling is a promising approach to offer enhanced traffic performance and also qualitative service differentiation without ubiquitous deployment of (1) direct NIAs that horizontally bind adjacent Network Planes, and (2) new routing protocols (e.g. q-BGP) based on which Parallel Internets are constructed. The motivation of this paradigm is based on the observation that the largest fraction of the traffic of an INP or its most important traffic is exchanged with only a few other INPs. In this case, it is possible to rely on a more lightweight approach relying on direct cooperation among these INPs to improve the distribution and performance of the traffic they exchange. The traffic exchanged between these cooperating INPs could traverse several domains that do not implement Network Planes and should be considered as BE domains. In spite of that it is possible to gain more control on the inter-domain paths followed by this traffic by deploying virtual peerings [QUOI05]. Virtual peerings are a combination of a management protocol and IP tunnelling that allows (1) to leverage the diversity of paths between two INPs, (2) to measure the performance of the obtained paths, and (3) to control which paths are used to actually carry the traffic. The IP tunnels are used to control through which peering link the traffic will enter a destination INP. Virtual peerings do not provide strict guarantees on the treatment of the traffic along the inter-domain paths since they may cross multiple BE domains. They can however support inter-domain traffic engineering actions such as load balancing by controlling through which peering links the traffic exits and enters the INP. However, it should be noted that where all INPs along the inter-domain paths engineered by the virtual peering mechanisms support service differentiation through NPs then virtual peerings may also be used to provision *qualitative* end-to-end service guarantees provided appropriate NIAs are in place.

7 FUNCTIONAL ARCHITECTURE

7.1 The overall architecture

Starting from the business model described in Section 3 this section analyses the interactions between the business roles of Customer, Service Provider and IP Network Provider, the functional blocks required to support those interactions and the internal functionality of INPs required to plan, engineer and operate Network Planes and Parallel Internets. Figure 31 depicts the interactions between the business roles. The arrows depict customer-provider relationships, with the arrowhead pointing towards from the provider to the customer and the labels identifying the type of agreement. It should be noted that SIAs and NIAs may also support peering relationships between two providers and this is modelled by two customer-provider relationships, one in each direction.

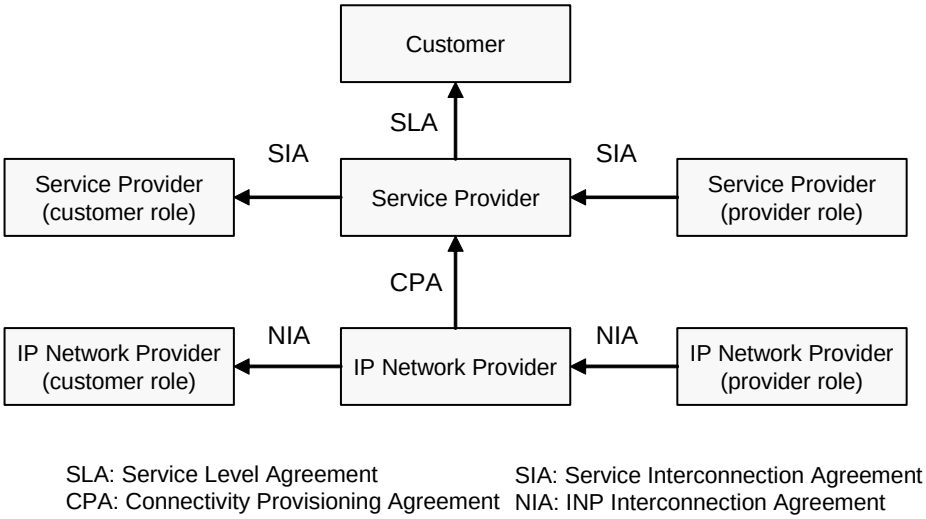

Figure 31 Interactions between Business Roles

Figure 32 presents the overall functional architecture from the perspective of the interfaces to be supported by each business role. In this diagram the directionality of the arrows indicate the direction of information flow. The upper three blocks of the SP form the customer interface and support interactions in relation to the Service Level Agreements (SLAs) between customers and SPs. The lower three blocks of the SP form the Connectivity Provisioning Agreement (CPA) interface with INPs. The three blocks on each side of the SP form the Service Interconnection Agreement (SIA) interface, while the corresponding blocks in the INP form the INP Interconnection Agreement (NIA) interface.

This interactions-focused view hides the complexity of internal functional blocks in the SP and INP which are contained within the Service Planning & Engineering and Network Plane Engineering meta-blocks, respectively.

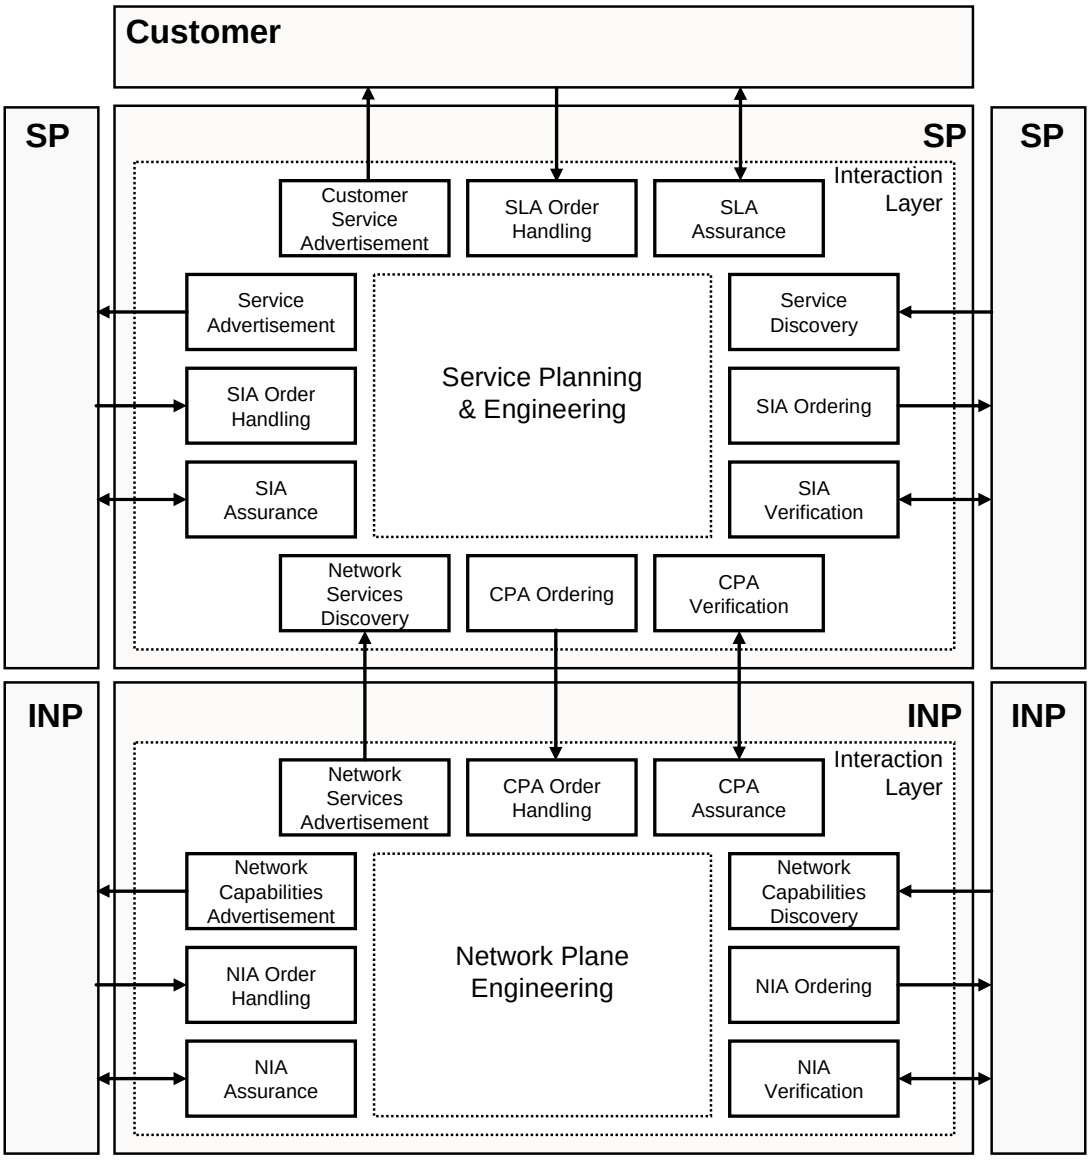

Figure 32 Overall AGAVE Functional Architecture: Interactions View

Figure 33 presents a more detailed decomposition of the INP, showing the functions involved with NP planning and engineering. The functional blocks depicted in both views are described at a high level in the following sub-sections.

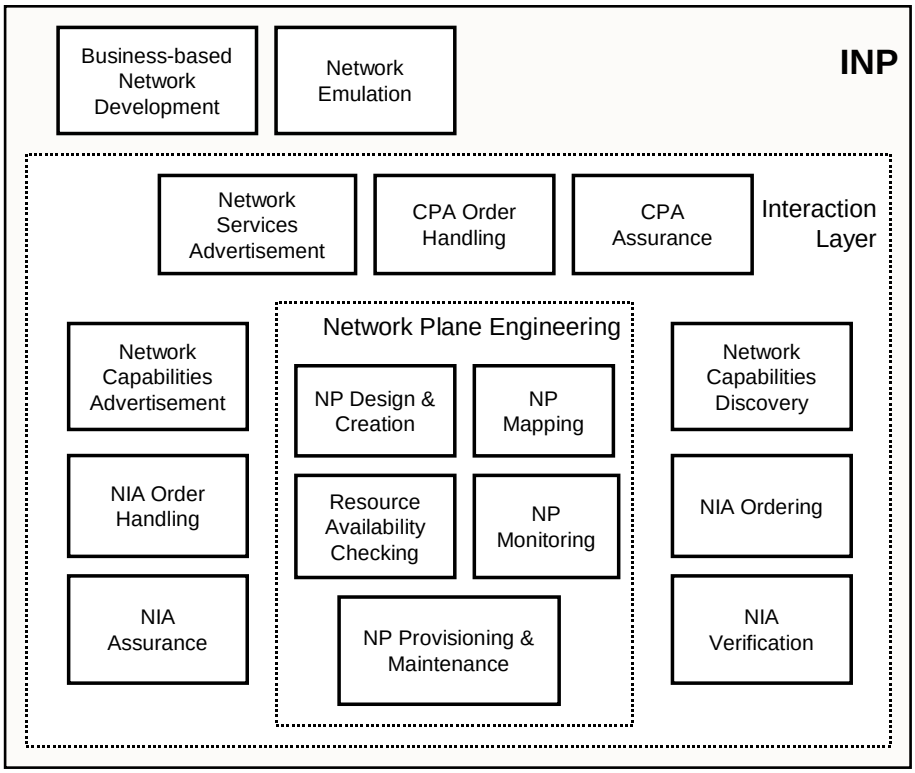

Figure 33 Detailed Functional Decomposition of the INP

7.2 Customer functional blocks

Using automatic or manual means, the customer executes well-defined service-specific procedures to discover the capabilities of the available service providers, to formulate and subscribe to a particular instance of the service, to make use of it, modify it, verify it and finally terminate it.

7.3 Service Provider functional blocks

7.3.1 Customer Interface

The *Customer Service Advertisement* function is responsible to announce the service provider's capabilities and service offerings, as generated by the *Service Planning and Engineering* function, to potential customers. *SLA Order Handling* function negotiates SLA establishment, modification or teardown requests with the customers. *SLA Assurance* function undertakes the interactions related to SLA assurance, involving scheduled or on-demand issuing of reports or event-triggered notifications concerning the compliance of the effective service to the SLA terms.

7.3.2 CPA Interface

The capabilities of the INPs are discovered by the *Network Services Discovery* function and fed to the *Service Planning and Engineering* function. The CPA order requests for establishment, execution of permissible actions, modification, explicit verification and teardown, formulated by the *Service Planning and Engineering* function, are negotiated with the INP through the *CPA Ordering* function. The processing of received CPA verification reports and event-triggered notifications sent by the INPs is undertaken by the *CPA Verification* function.

7.3.3 SIA Interface

Service Planning and Engineering announces the available services of the service provider through the *Service Advertisement* function and discovers the corresponding capabilities of other service providers

through the *Service Discovery* function. The *SIA Order Handling* function receives SIA orders for establishment, execution of permissible actions, modification or teardown issued by peer service providers. The actual implementation of the request is undertaken by the *Service Planning and Engineering* function. The SIA order requests generated by the *Service Planning and Engineering* function are negotiated with the implicated peer service providers through the *SIA Ordering* function. The SIA assurance reports and notifications generated by the *Service Planning and Engineering* function are sent through the *SIA Assurance* function. The processing of received SIA verification reports and event-triggered notifications sent by peer service providers is undertaken by the *SIA Verification* function.

7.3.4 Service Planning & Engineering

The *Service Planning and Engineering* functional block is the place where the offered services are designed, created, implemented, maintained, packaged and marketed. This block is responsible for identifying the candidate SPs to extend the scope of its services and the candidate INPs to ensure its connectivity requirements. It is also responsible to set guidelines for the Interaction Layer to initiate, modify or terminate SIAs with other service providers and CPAs with INPs.

7.4 IP Network Provider functional blocks

7.4.1 CPA Interface

The *Network Services Advertisement* function is responsible to announce the IP network provider's connectivity provisioning capabilities, as formulated by the *NP Engineering* functions (see below), to potentially interested service providers. *CPA Order Handling* function negotiates CPA establishment, modification or teardown requests with the service providers. The actual CPA fulfilment and the execution of control actions invoked by the service providers (permissible within the context of their CPAs) are undertaken by the *NP Engineering* functions. The *CPA Assurance* function undertakes the interactions related to CPA assurance, involving scheduled or on-demand issuing of reports or event-triggered notifications concerning the compliance of the effective connectivity provisioning to the CPA terms.

7.4.2 NIA Interface

The *NP Design and Creation* function announces the interconnection capabilities of the IP network provider through the *Network Capabilities Advertisement* function. The corresponding capabilities of other IP network providers are discovered through the *Network Capabilities Discovery* function and fed to the *NP Engineering* functions. The *NIA Order Handling* function receives NIA order requests for establishment, execution of permissible actions, modification or teardown issued by peer IP network providers. The actual implementation of the request is undertaken by the *NP Engineering* functions. The NIA order requests generated by the *NP Engineering* functions are negotiated with the implicated peer IP network providers through the *NIA Ordering* function. The NIA assurance reports and notifications generated by *NP Engineering* functions are sent through the *NIA Assurance* function. The processing of received NIA verification reports and event-triggered notifications sent by peer IP network providers is undertaken by the *NIA Verification* function.

7.4.3 Network Plane Planning & Engineering

Network Plane Planning and Engineering is composed of two functional block groups as illustrated in the figure below.

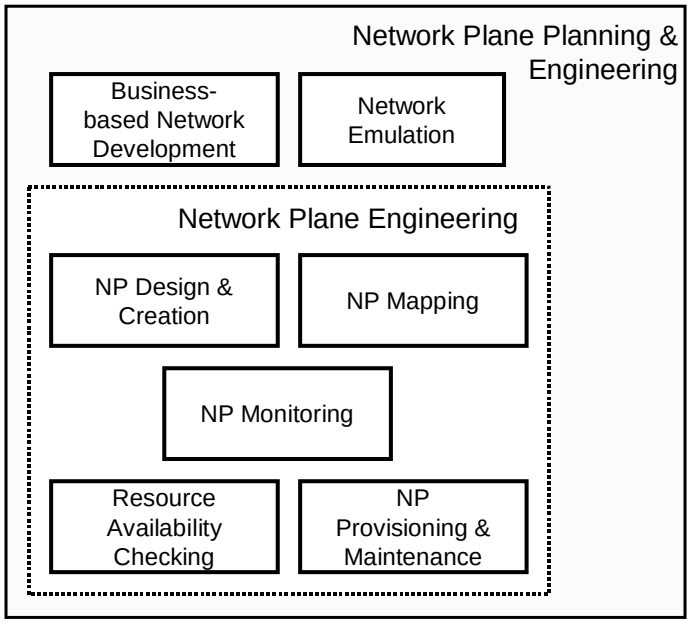

Figure 34 Network Plane Planning & Engineering Block

The first group, which includes *Business-based Network Development* and *Network Emulation* functional blocks, is responsible for the planning of network operations, production of evolution roadmaps and strategy of the network, expansion decisions of the network capabilities and acceptance of service provisioning requests received from service providers. This group interacts with the second group, called *Network Plane Engineering*, in order to map its requirements into engineering tasks and network operations. More details about individual components of *Network Plane Engineering* are provided in the following sections.

7.4.3.1 Business-based Network Development

The *Business-based Network Development* functional block is responsible for the development of the network capabilities provided to other INPs and the IP connectivity services offered to Service Providers. It is aware of the peering business opportunities thanks to the input received from *Network Capabilities Discovery*. It is also aware of the requirements of the service providers and the profit that can be achieved if some CPA agreements are established.

7.4.3.2 NP Emulation

The ultimate goal of the *NP Emulation* functional block is to provide the *Business-based Network Development* functional block with data to support its decision-making process, based on its knowledge of the local capacities and environment, regarding the impact (financial, engineering, service capabilities, etc) on currently engineered Network Planes, offered IP connectivity provisioning capabilities, etc.

7.4.3.3 NP Engineering

NP Engineering functional block is responsible for the design, activation, provisioning and maintenance of Network Planes. This functional block is composed of five sub-functional blocks as described below:

7.4.3.3.1 NP Design & Creation

- The *NP Design & Creation* functional block is responsible for the off-line Network Planes dimensioning before actual enforcement into operational networks of a given INP. The design and

creation phase produces high level specifications of the Network Planes in terms of qualitative and quantitative parameters associated with each dimension as defined in section 6.2.2.

7.4.3.3.2 NP Provisioning & Maintenance

- *NP Provisioning and Maintenance* functional block is responsible for operations on the network and actual enforcement of configuration tasks on network nodes and functions. This functional block is the place where the *NP Design and Creation* output is translated into operational tasks.

7.4.3.3.3 NP Monitoring

The *NP Monitoring* functional block is responsible for monitoring the activated Network Planes within the network of an INP. NP monitoring includes monitoring activities per NP, CPA and NIA. It interacts with: *CPA/NIA Assurance* to provide appropriate data as agreed in the CPA/NIA assurance clauses; *NP Provisioning and Maintenance* to provide information to a dynamic online network provisioning process; *NP Design and Creation* to send notifications on critical performance deterioration events; and finally with *NP Emulation* to adjust the traffic profiles it uses.

7.4.3.3.4 NP Mapping

NP Mapping functional block is responsible for assessing the CPA/NIA requirements in terms of traffic demand with particular QoS and resilience guarantees and for finding the best NP candidate(s) to accommodate these requirements.

7.4.3.3.5 Resource Availability Checking

- *Resource Availability Checking* functional block is responsible for tracking the current status of the network resources. It keeps track of allocated resources per Service Provider, per CPA, per peer INP and per NIA.

7.5 Conclusion

This section has presented a high-level description[2] of the AGAVE functional architecture. The detailed design of functional blocks and their interactions will continue to be developed in the course of the detailed specification tasks in WP2 and WP3.

[2] More detailed specifications of the functional blocks and interaction scenarios describing information flows between them have been created by the project and these are currently documented in an internal version of this deliverable. A full set of specifications will be made public at a later stage.

8 SUMMARY

This document addresses the problem of service provisioning and delivery across the Internet from the standpoints of service provider and IP network provider. An interface is defined to enable their smooth interaction for the provisioning of the IP connectivity required by the IP-based services. As a means to accommodating services with diverse IP connectivity provisioning requirements running over a common IP infrastructure, the document outlines the Parallel Internets framework and architecture.

Specifically, a *Business Model* is drawn to capture the business roles and relationships pertinent to the end-to-end deployment of IP-based QoS-enabled services. A clear distinction between the roles of Service Provider and IP Network Provider is adopted, creating a business opportunity for specialised service providers to control without owning the IP infrastructure, and a new stream of revenue for the IP network providers who actually own it. Acknowledging the requirement for global coverage at both the IP and the service layers, providers are associated by horizontal interconnection relationships.

IP connectivity, Voice over IP and Virtual Private Network services are studied as *Business Cases* to set the requirements for the Parallel Internets framework and to drive the AGAVE work. Current business practices, latest developments, desirable enhancements and associated issues are investigated for each of the selected services. The *Requirements* of customers and service providers specific to each selected business case are envisaged.

A *High-level Specification* of the *Interactions* among service actors is provided focusing on the IP connectivity provisioning interactions, necessary to fulfil the derived requirements. IP network providers interact with each other to expand the scope of their offerings. Service providers interact with IP network providers to control the provisioning of the IP connectivity portion assigned to their services. The interactions at the service layer are elaborated for each selected business case to demonstrate and validate the IP connectivity provisioning interactions.

The *Network Plane* concept is introduced and aspects of the Network Plane engineering process are investigated. Network Planes are brought into play for accommodating traffic with different end-to-end QoS and resilience requirements, coming from different service providers operating different service types. Network Planes are interconnected to build *Parallel Internets*. Alternative techniques for Network Plane and Parallel Internets realisation are presented.

A *Functional Architecture* is specified, covering the functionality required at both the service and the IP layer for the deployment of end-to-end QoS-enabled services following the paradigm of the Parallel Internets. Vertical and horizontal interfaces are defined between service actors. Service and IP layer functional blocks are identified and outlined.

This document constitutes the formal output of WP1, set with the objective to form the framework for the work in WP2 and WP3. The results of WP1 work are not conclusive; rather they set the principles and the starting point for the specification work to follow.

Based on the requirements of the considered business cases and on the capabilities provided by Network Planes and Parallel Internets, WP2 undertakes the detailed specification, design and development of the IP connectivity provisioning interface between the service and IP network providers The starting point of the WP2 work will be the high-level specification of interactions and interaction layer functional blocks presented in this document.

WP3 undertakes the detailed specification, design and development of appropriate mechanisms and algorithms for realising the Network Planes and Parallel Internets. To this end, WP3 will be initially based and appropriately revise the high-level specifications captured in this document.

The detailed specification, design and implementation work in WP2 and WP3 will validate the feasibility of the approach outlined in this document.

9 REFERENCES

[ACTRICE] French RNRT (Réseau National de Recherche en Télécommunication) research projects, for more information visit ACTRICE URL: http://eurongi.enst.fr/actrice/news.php

[ALME99a] Almes, G., Kalidindi, S., Zekauskas, M., *A One-way Delay Metric for IPPM*, RFC2679, September 1999

[ALME99b] Almes, G., Kalidindi, S., Zekauskas, M., *A One-way Packet Loss Metric for IPPM*, RFC2680, September 1999

[ALME99c] Almes, G., Kalidindi, S., Zekauskas, M., A Round-trip Delay Metric for IPPM, RFC2680, September 1999

[AWDU02] Awduche, D., Chiu, A., Elwalid, A., Widjaja, I., Xiao, X., *Overview and Principles of Internet Traffic Engineering*, RFC3272, May 2002

[BARL75a] Barlow, R.E., Proschan, F., Statistical theory of reliability and life testing. New York: Holt, Rinehart and Winston, Inc., 1975

[BARL75b] Barlow, R.E., Proschan, F., (1975), *Statistical Theory of Reliability and LifeTesting: Probability Models*, Holt, Rinehart and Winston, Inc

[BATE02] Bates, T. et al, *Multiprotocol Extensions for BGP-4*, RFC 2858

[BAUG04] Baugher, M. et al., *The Secure Real-time Transport Protocol (SRTP)*, RFC3711, March 2004.

[BANG05] Bangalore, M. et al, *A Telephony Gateway REgistration Protocol (TGREP)*, IETF Internet Draft, draft-ietf-iptel-tgrep-06.txt, July 2005

[BERN98] Berners-Lee, T. et al., Uniform Resource Identifiers (URI): Generic Syntax, RFC2396, August 1998

[BIER02] Bierman, A., *Remote Monitoring MIB Extensions for Differentiated Services*, RFC3287, July 2002

[BOUC05] Boucadair, M., *QoS-Enhanced Border Gateway Protocol*, IETF Internet-Draft, draft-boucadair-qos-bgp-spec-01.txt, July 2005

[BRYA05] Bryant, S., Pate, P., *Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture*, RFC3985, March 2005

[CAMA04] Camarillo, G. et al, *Integration of Resource Management and Session Initiation Protocol (SIP)*, RFC3312, October 2002

[CAMA04] Camarillo, G. et al, *Integration of Resource Management and Session Initiation Protocol (SIP)*, RFC3312, October 2002

[DEMI02] Demichelis, C., Chimento, P., *IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)*, RFC3393, November 2002

[ENNS06] Enns, R., *NETCONF Configuration Protocol*, IETF Internet-Draft, draft-ietf-netconf-prot-12, February 2006

[EUQOS] European IST (Information Society Technologies) research projects, for more information visit IST-EUQOS URL: http://www.euqos.org

[FALT04] Faltstrom, P., Mealling, M., *The E.164 to Uniform Resource Identifiers (URI) Dynamic Delegation Discovery System (DDDS) Application (ENUM)*, RFC3761, April 2004

[FARI00] Farinacci, D. et al, *Generic Routing Encapsulation (GRE)*, RFC 2784, March 2000

[FEAM03] Feamster, N. et al, *Guidelines for interdomain traffic engineering*, ACM SIGCOMM Computer Communications Review, 33(5), 2003, pp. 19-30

[FORT02] Fortz, B. et al, *Traffic engineering with traditional IP routing protocols*, IEEE Communication Magazine, Vol. 40, Issue 10, 2002, pp. 118-124

[HALS06] Halstead, M. et al., *Requirements for Multi Autonomous System VPN Services*, Internet draft, draft-halstead-guichard-mavs-requirements-02, May 2006

[HABE06] Haberler, M., Hammer, M., Lendl, O., *A Federation based VoIP Peering Architecture*, IETF Internet Draft, draft-lendl-speerimint-federations-01, June 2006

[HUST04] Huston, G., *NOPEER Community for Border Gateway Protocol (BGP) Route Scope Control*, RFC3765, April 2004

[INTERMON] European IST (Information Society Technologies) research projects, for more information visit IST-INTERMON URL: http://www.ist-intermon.org

[IPSF] www.ipsphere.org

[KENT98a] Kent, S., Atkinson, R., *IP Authentication Header*, RFC 2402, November 1998

[KENT98b] Kent, S., Atkinson, R., *IP Encapsulating Security Payload (ESP)*, RFC 2406, November 1998

[KOMP05] Kompella, K., Rekhter, Y., *Virtual Private LAN Service*, draft-ietf-l2vpn-vpls-bgp-06, December 2005

[KOOD02] Koodli, R., Ravikanth, R., *One-way Loss Pattern Sample Metrics*, RFC3357, August 2002

[KULM06] Kulmala, M. et al, *ASBR VRF context for BGP/MPLS IP VPN*, draft-kulmala-l3vpn-interas-option-d-02, February 2006

[LAU05] Lau, J., Townsley, M., Goyret, I., Kent, S., Atkinson, R, *Layer Two Tunneling Protocol - Version 3 (L2TPv3)*, RFC 3931, March 2005

[LASS05] Lasserre, M., Kompella, Y., *Virtual Private LAN Service over MPLS*, IETF Internet Draft, draft-ietf-l2vpn-vpls-ldp-08, November 2005

[LOBSTER] European IST (Information Society Technologies) research projects, for more information visit IST-LOBSTER URL: http://www.ist-lobster.org/about/objectives.html

[MULE05] Mule, JF., *SPEERMINT Requirements for SIP-based VoIP Interconnection*, IETF Internet Draft, draft-ietf-speermint-requirements-00.txt, June 2006

[MAHD99] Mahdavi, J., Paxson, V., *IPPM Metrics for Measuring Connectivity*, RFC2678, September 1999

[MESCAL] European IST (Information Society Technologies) research projects, for more information visit IST-MESCAL URL: http://www.ist-mescal.org

[MOME] European IST (Information Society Technologies) research projects, for more information visit IST-MOME URL: http://www.ist-mome.org

[NAGA04] Nagarajan, A., *Generic Requirements for Provider Provisioned Virtual Private Networks (PPVPN)*, RFC3809, June 2004

[NOBEL] European IST (Information Society Technologies) research projects, for more information visit IST-NOBEL URL: http://www.ist-nobel.org

[ONOK05a] Ono, K., Tachimoto, S., *Requirements for End-to-Middle Security for the Session Initiation Protocol (SIP)*, RFC 4189, October 2005.

[ONOK05b] Ono, K., Tachimoto, S., *End-to-middle Security in the Session Initiation Protocol (SIP)*, IETF Internet Draft, draft-ietf-sip-e2m-sec-01, October 2005

[PATR01] Patrick, M., *DHCP Relay Agent Information Option*, RFC3046, January 2001

[PRZY05] Przygienda, T. et al, *M-ISIS: Multi Topology (MT) Routing in IS-IS*, IETF Internet Draft, draft-ietf-isis-wg-multi-topology-09.txt, March. 2005, work in progress

[PSEN05] Psenak, P. et al, *Multi-Topology (MT) Routing in OSPF*, IETF Internet Draft, draft-ietf-ospf-mt-04.txt Apr. 2005

[PERK96] Perkins, C., *IP Encapsulation within IP*, RFC 2003, October 1996

[QUOI03] Quoitin, B. et al, *Interdomain traffic engineering with BGP*, IEEE Communications Magazine, May 2003

[QUOI05] B. Quoitin and O. Bonaventure, "A cooperative approach to interdomain traffic engineering", Proc. 1st Conference on Next Generation Internet Networks Traffic Engineering (NGI) 2005

[ROSE02] Rosenberg, J. et al., *SIP: Session Initiation Protocol*, RFC 3261, June 2002.

[ROSE99] Rosen, E., Rekhter, Y., *BGP/MPLS VPNs*, RFC 2547, March 1999

[ROSE00] Rosenberg, J., Schulzrinne, H., *A Framework for Telephony Routing over IP*, RFC 2871, June 2000

[ROSE06] Rosen, E., Rekhter, Y., *BGP/MPLS IP Virtual Private Networks (VPNs)*, RFC 4364, February 2006

[RAIS02] Raisanen, V., Grotefeld, G., Morton, A., *Network performance measurement with periodic streams*, IETF, November 2002

[RIPE06] RIPE weekly routing table report http://www.ripe.net/ripe/maillists/archives/routing-wg/2006/msg00103.html

[ROWS01] Rowstron, A., Druschel, P., *Pastry: Scalable, Decentralized Object Location, and Routing for Large-Scale Peer-to-Peer Systems*, in proceedings of Middleware 2001: IFIP/ACM International Conference on Distributed Systems Platforms, Heidelberg, Germany, November 12-16, 2001

[SALS02] Salsano, S., Veltri, L., *QoS Control by means of COPS to support SIP based applications*, IEEE Network, March/April 2002

[SCAMPI] European IST (Information Society Technologies) research projects, for more information visit IST-SCAMPI URL: http://www.ist-scampi.org

[SCHW06] Schwartz, D., et al., *SPAM for Internet Telephony (SPIT) Prevention using the Security Assertion Markup Language (SAML)*, IETF Internet Draft, draft-shwartz-sipping-spit-saml-01, June 2006

[SRID03] Sridharan, A. et al, *Achieving Near-Optimal Traffic Engineering Solutions for Current OSPF/IS-IS Networks*, Proc. IEEE INFOCOM, pp. 1167-1177, Apr. 2003

[STOI01] Stoica, I., Morris, R., Karger, D., Kaashoek, M.F., Balakrishnan H., *Chord: A scalable peer-to-peer lookup service for internet applications*, in proceedings of SIGCOMM 2001, August 2001

[TEQUILA] European IST (Information Society Technologies) research projects, for more information visit IST-TEQUILA URL: http://www.ist-tequila.org

[TISPAN] http://portal.etsi.org

[VELT02] Veltri, L., Salsano, S., Papalilo, D., *SIP Extensions for QoS support*, IETF Internet Draft, draft-veltri-sip-qsip-01, September 2002

[VELT03] Veltri, L., Salsano, S., Papalilo, D., *QoS Support for SIP Based Applications in DiffServ Networks*, In Proc. SoftCom2003

[WALD00] Waldbusser, S., *Remote Network Monitoring Management Information Base*, RFC2819, May 2000

[WALD04] Waldbusser, S., *Application Performance Measurement MIB*, RFC3729, March 2004

[WALD99] Waldbusser, S., *Remote Network Monitoring Management Information Base Version 2 using SMIv2*, RFC2021, January 1999

[Wu05] Wu, T., *Network Neutrality*, Broadband Discrimination, Journal of Telecommunications and High Technology Law, Vol. 2, p. 141, Available at SSRN: http://ssrn.com/abstract=388863 or DOI: 10.2139/ssrn.388863, April 2005

[Wu04] Wu, T., *The Broadband Debate: A User's Guide*, Journal of Telecommunications and High Technology Law, Vol. 3, No. 69, Available at SSRN: http://ssrn.com/abstract=557330, 2004

[Y.1540] *IP packet transfer and availability performance parameters*, ITU-T Recommendation, December 2002

[Y.1541] *Network performance objectives for IP-based services*, ITU-T Recommendation, May 2002

[6QM] European IST (Information Society Technologies) research projects, for more information visit IST-6QM URL: http://www.6qm.org/index.php

[3GPP] http://www.3gpp.org/

10 ABBREVIATIONS

AGAVE	A liGhtweight Approach for Viable End-to-end IP-based QoS Services
AS	Autonomous System
BGP	Border Gateway Protocol
CE	Customer Edge
CPA	Connectivity Provisioning Agreement
CPE	Customer Premise Equipment
DDoS	Distributed Denial of Service
DIFFSERV	Differentiated Services
DNS	Domain Name System
DoS	Denial of Service
DSCP	Differentiated Services Code Point
FIB	Forwarding Information Base
GW	Gateway
INP	IP Network Provider
ISP	Internet Service Provider
ITAD	IP Telephony Domain
L1VPN	Layer-1 VPN
L2VPN	Layer-2 VPN
L3VPN	Layer-3 VPN
LI	Legal Interception
LS	Location Server
LSP	Label Switched Path
MPLS	Multi Protocol Label Switching
NIA	Network Interconnection Agreement
NP	Network Plane
OSPF	Open Shortest Path First
PDB	Per-Domain Behaviour
PHB	Per-Hop Behaviour
PI	Parallel Internet
PS	Proxy Server
QC	QoS Class
QoS	Quality of Service
RFC	Request For Comment
RIB	Routing Information Base
RIP	Routing Information Protocol
SBC	Session Border Controller
SIA	Service Interconnection Agreement
SIP	Session Initiation Protocol
SLA	Service Level Agreement
SLS	Service Level Specification
SP	Service Provider

11 APPENDIX A: ISSUES IN IP QOS

11.1 Intrinsic QoS versus perceived QoS

The basic ITU definition of Quality of Service (QoS) expressed for the first time in E.800 is: "*the collective effect of service performance which determines the degree of satisfaction of a user of the service*". Following this definition, the provisioning of QoS should be as simple as aligning the service performance as perceived by the end user with the end user's expectations. The problem of this approach is that the perception of the end user has a key role and this cannot be directly controlled by the Service Provider, as this perception involves great doses of subjectivity. Therefore, it is convenient to distinguish between intrinsic QoS and perceived QoS:

- The Intrinsic QoS pertains to service features arising from technical aspects and it is determined by the network design and its provisioning. The required quality is achieved through an appropriate network design, a selection of transport protocols, Quality of Service assurance mechanisms and selection of related parameter values. Intrinsic QoS parameters are defined at network level. Additional parameters at the service level should also be defined. The user perception of the service does not influence in the intrinsic QoS rating.
- The Perceived QoS reflects the user experience while using a particular service. It is influenced by the user expectations compared to the actual observed service performance. In turn, the personal expectations are usually affected by the user previous experiences with similar telecommunications services and other user beliefs. Thus, the quality of a service with given intrinsic features may be perceived differently by various users.

When engineering NPs and PIs, INPs are concerned with network level performance and how this supports the service requirements as agreed in CPAs with SPs. Therefore, according to the above definitions, INPs are only concerned with intrinsic QoS. Perceived QoS is an issue to be managed by SPs through their relationships with customers and this is outside the scope of AGAVE.

11.2 QoS *vs.* "Network Neutrality"

The "Network Neutrality" topic is currently being discussed in the USA, thus it seems necessary to clearly establish a position regarding this issue as not only AGAVE deals with QoS, which could be considered as hurdle to "Network Neutrality" but also this discussion is bound to arise in Europe in the near future. "Network Neutrality" defenders advocate that networks should be as neutral as possible so that services deployed over them can compete in fair conditions (i.e., without discriminating any service or traffic belonging to a given Service Provider) and hence allowing for maximising innovation and welfare. To achieve this, and in order to prevent potential discriminating activities by IP Network Providers (while trying to expand their market boundaries by integrating vertically in the value chain may be tempted to discriminate third Service Providers), their defenders ask to the government to intervene so that the "Network Neutrality" be fostered and preserved.

Carrying the "Network Neutrality" concept to extremes, Quality of Services mechanisms would not be considered as neutral, as they produce some kind of discrimination due to the fact that some services are treated better than others (*even if it is arguable that "Network Neutrality" paradigm apply inside the same QoS class and not to all QoS classes. Another question to answer is whether activating QoS mechanisms will impact the performance of BE packets and evaluate this impact*). In this way, "Network Neutrality" extremists advocate for preserving the "best-effort" Internet, where the network makes the same effort to carry each packet.

Notwithstanding, and as Tim Wu, a defender of "Network Neutrality", says in [Wu05], native IP technology is somehow biased against real-time applications. This is due to the aim that the original Internet was designed to fulfil military survivability, so that it was supposed to be unreliable in nature, meaning that the applications running on either end had to be prepared to recognize data loss and retransmitting data as many times as necessary to achieve its ultimate delivery. This principle of

unreliable delivery means that the Internet only makes a "best-effort" attempt to deliver packets, the network can drop a packet without any notification to the sender and the receiver.

On the other end, the deregulationists' position consists of asking for a minimal regulation framework, or even no regulation at all (i.e., laissez faire), so that there are enough incentives to go on investing on, rather expensive, telecommunication networks. According to [Wu04], these two positions are not necessarily irreconcilable. Indeed, Michael Powell, FCC chairman, defends the normative desirability of "Internet freedom" by "*ensuring that consumers can obtain and use the content, applications and devices they want*". Powell's discussion of "Internet freedom" focuses on users' rights. The four freedoms are:

- Freedom to Access Content. Consumers should have their choice of legal content;
- Freedom to Use Applications. Consumers should be able to run applications of their choices unless they exceed service plan limitations or harm the provider's network;
- Freedom to Attach Personal Devices. Consumers should be permitted to attach personal devices they choose to the connections that they pay for in their homes so long as the devices operate within service plan limitations and do not harm the provider's network or enable theft of service;
- Freedom to Obtain Service Plan Information. Consumers must receive clear and meaningful information regarding their service plans and what the limits of those plans are.

Obviously, these freedoms are not biased against the implementation of QoS mechanisms by the network operator. So, it is licit to offer connectivity services with QoS, provided that fair play is respected (i.e., every service provider may use these QoS capabilities under the same conditions or the users get to choose the QoS they want to use), and that conditions imposed by the INP be transparent to the user so that he knows what he is paying for.

Furthermore, it is obvious that the concept of "Network Neutrality" is aligned with the spirit of IP and that is suitable in some business models. But, in order for the Internet to live, the business of the INPs should also be maintained since they are the ones who build and maintain IP infrastructures.

12 APPENDIX B: TERMINOLOGY FOR RESILIENCE IN IP NETWORKS

This appendix provides a list of useful definition so as to reflect the notion of resilience, mainly within IP networks. These definitions can be applied to IP service in general. Some examples and literature pointers have been furnished so as to ease the understanding of these notions.

- The MTBF measures the average time spend in the "up" state between two successive failures.
- The MTTR denotes the "Mean Time To Repair" or the "Mean Time To Restore".
    - The "Mean Time To Repair" includes all the actions actually necessary to restore the (IP) service: failure detection, provisioning of the spare components, the repair itself, restarting (making the necessary updates, reboot the machine, etc.)
    - The "Mean Time To Restore", in a IP/MPLS network which have an alternate path, measures the rerouting time. It includes the failure detection time, the failure advertisement, the computation of an alternate path, the installation of theses alternates' paths in the forwarding tables. The whole process may require multiple steps to find and select the best alternate route.
- The availability is well established in the literature, for instance [BARL75a] or [BARL75b] which define the availability of a repairable system as "the probability that the system is operating at a specified time t". In telecommunications and reliability theory, the term availability has the following meanings:
    - The degree to which a system or equipment is operable and in a committable state at the start of a mission, when the mission is called for at an unknown, i.e., a random, time. Simply put, availability is the proportion of time a system is in a functioning condition.
    - The ratio of (a) the total time a functional unit is capable of being used during a given interval to (b) the length of the interval.
        - An example of availability is 100/168 if the unit is capable of being used for 100 hours in a week.
        - Typical availability objectives are specified either in decimal fractions, such as 0.9998, or sometimes in a logarithmic unit called nines, which corresponds roughly to a number of nines following the decimal point, such as "five nines (5 9')" for 0.99999 reliability.
    - The availability can easily be computed knowing the MTTR and the MTBF: availability = MTBF/(MTTR+MTBF).
- The maintenance window is a period of time -usually off business hour- when INP and/or SP perform their planned service operations. Each INP/SP may have different maintenance windows.

13 APPENDIX C: OVERVIEW OF IP ROUTING ISSUES TO SUPPORT NP ENGINEERING

There are several ways of classifying routing mechanisms which involve different layers and strategies (e.g., hop by hop routing versus explicit routing, static versus dynamic routing or single-path versus multi-path routing, single-layer versus multi-layer, etc.). However, the AGAVE project will focus on IP routing layer leaving the complementary and alternative routing management strategies of both optical and link layer out of the scope of the current project. Nevertheless, further information can be found in NOBEL IST project [NOBEL].

One of the main challenges of current IP networks is how to carry a wide range of traffic types while complying with their different QoS performance requirements and making an efficient use of network resources, particularly the available bandwidth. Congestion situations usually arise when the network resources are inadequate or insufficient to carry the offered load, or when traffic flows are inadequately allocated to available resources. Route management is a key element of Traffic Engineering, because efficient management and control of network resources while providing QoS can be achieved only by the use of an adequate routing strategy.

Routing is the process of discovering and selecting paths in a network to carry the traffic from a given source to every destination. Given the different QoS requirements of the traffic patterns carried on these networks and the uncertainty of traffic demand, it is frequently necessary to use specific policies to distribute load between the available resources in order to avoid congestion states. Hence, an adequate design of the protocol that discovers and selects the path or route that a given traffic flow follows is extremely important to achieve the best possible results in terms of performance.

On the one hand, routing strategies can be classified in function of how the routing is solved: hop-by-hop or explicit routing mechanisms (e.g.: MPLS):

- In hop-by-hop based routing, each node works autonomously choosing the next hop to reach a given destination. This kind of mechanisms could either have a partial view of the network (distant vector algorithms) or a full view (link-state algorithms). In the latter case, it is required that each node should have consistent view on the overall network topology in order to avoid loops and slow convergence problems.
- Explicit routing algorithms are usually centralised as end-to-end paths are established by means of a central control system. Nevertheless, paths can also be established in a distributed way (e.g., LSPs made based on IGP routing).This kind of mechanisms normally provides a tight control of network paths, but can present scalability issues.

On the other hand, routing can be classified into inter and intra domain. The next section will deal with this last classification.

13.1.1.1 Inter-domain & Intra-domain Routing

IP routing can be also classified into two groups: intra-domain and inter-domain ones. The first type of routing paradigms are used to find the path inside an administrative domain (from now on, we will assume that each domain is an Autonomous System -AS-), whereas the second ones are used to find the path across multiple domains (ASes).

IP networks typically use hop-by-hop protocols for intra-domain routing —Routing Information Protocol (RIP), Open Shortest Path First (OSPF), Intermediate System to Intermediate System (IS-IS)—, which take routing decisions locally, according to simple algorithms of path computation. In these protocols the path computation is commonly based on static metrics such as hop-count and link weights. When using link state based protocols such as OSPF and IS-IS, network administrators are often recommended to configure link weights according to static parameters associated with network links, such as delay and bandwidth capacity. This simple approach allows an efficient deployment of IP networks.

However, these path selection algorithms may introduce problems related to convergence and control of traffic flows. This situation has led to proposals that try to overcome some of the disadvantages of intra-domain IP routing, without losing its advantages. The objective of AGAVE is to deploy mechanisms to improve network performance and facilitate the construction of Network Planes without introducing an excess of complexity.

Inter-domain routing in the Internet takes place between different ASes. Nowadays, Border Gateway Protocol (BGP) is the de facto inter-domain routing protocol and it is used to exchange routing information among different administrative domains. However, the exchanged information does not allow path selection based on QoS constrains not allowing the realisation of Parallel Internets. Due to these limitations, q-BGP [BOUC05] has been developed as a means to support inter-domain QoS requirements. In addition to these improvement proposals for IP routing strategies, MPLS can be used as an alternative to IP routing allowing the creation of Network Planes and the establishment of end-to-end paths with given constraints.

13.1.1.2 Differentiated Routing

Routing can be classified on the basis of the means to find the path to reach a destination from a given origin. For instance, the whole path can be calculated by the origin (as in virtual circuit technologies such as ATM or MPLS). Another option consists of selecting the best path hop by hop (as IP routing protocols do). A third option is flooding, so that each packet be delivered to every possible destination (as Ethernet flooding mechanism does). This last option is a good choice for delivering broadcast contents.

So far, most routing mechanisms have been developed with the aim of using network resources in the most efficient way; whereas the treatment of different QoS-sensitive traffic has been generally assigned to forwarding mechanisms, such as DiffServ. Meaning that all the traffic carried from an origin to a destination follows the same path, no matter its QoS constrains, relying on the forwarding policies to provide the required differentiation.

However, lately some routing protocols have been designed to take QoS into account (e.g. q-BGP). This paradigm aims to design routing protocols that are able to provide a differentiated treat (path) to different services is gaining importance. In addition, MPLS has also been extended so as to be used as a mechanism of explicit routing for supporting QoS. Specifically, label switched paths can be constructed for achieving the demanded QoS requirements.

If these protocols are to offer different paths to different types of traffic they need to be based on techniques that enable routing with multiple paths. This means that different paths must be used to carry traffic from a given source to a destination and an optimal path chosen according to the QoS constrains (delay, jitter, packet loss, etc.) of a given type of traffic. The term ‘optimal path’ here refers to the path that satisfies the QoS requirements and is also efficient in network resource utilisation.

It has to be noticed that for standard IP routing protocols, normally the optimum path is that with minimum cost (calculated by number of hops, delay or distance), without taking into account the load in the network. Individually, every flow would like to go through this path. However, taking into account the network load, some flows are routed by ‘sub-optimal’ paths with higher cost but improving the overall network performance.

MPLS is a good example of this kind of routing strategies because a given flow can be assigned to different LSPs and thus to different paths depending on its QoS requirements. Nevertheless, there are other types of multi-path routing algorithms that also take into account QoS traffic classes, so that each QoS Class is routed in a fashion that guarantees required QoS. For instance, a strategy where low priority traffic was diverted to the most costly paths, would reserve the optimal paths (those with lower costs) to the high priority traffic and thus could help to provide the required QoS.

14 APPENDIX D: INTER-AS VPN STATE OF THE ART

A VPN can be defined as a set of transmission and switching resources (network) dedicated to a customer (private) and built over a shared infrastructure (virtual).

A L3 VPN processes (i.e. forwards) the traffic at the IP layer and provides routing. A L2 VPN processes (i.e. switches) the traffic at the layer 2 (Frame Relay DLCI, ATM VP/VC, Ethernet VLAN, Ethernet MAC address).

14.1 VPN Taxonomy

The taxonomy of existing VPNs is presented in the following sections.

14.1.1 Overlay model

In this model, CEs directly peers with each others as follows:

- The IP network provider only provides a plain layer 2 or 3 connectivity and is VPN unaware;
- The CEs create the VPN service on top (overlay) of the providers network by setting tunnels among themselves and running any required protocol inside the tunnels (IP IGP routing protocols for a L3 VPN).

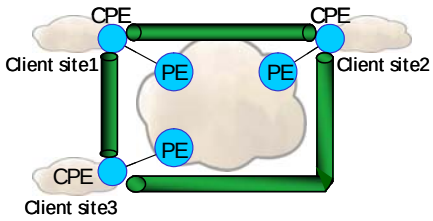

CPE-VPN Model

An overlay VPN can be built over Layer 2 networks (ATM, FR, etc) or Layer 3 with the use of IP Tunnels (IP-in-IP [PERK96], GRE [FARI00], IPSec [KENT98a], L2TP [LAU05]).

This overlay VPN is usually built by the Customer itself, but may be outsourced to a Service Provider.

14.1.2 Peer model

In the peer model, the customer network edge node -the CE equipment- peers with the VPN Service Provider's Edge (PE) equipment. CEs do not peer with each others.

To provide traffic isolation between VPNs and the shared public network, the VPN Service Provider uses tunnels between its PE nodes. Typically, MPLS tunnels are used, but one could also use IP tunnels (GRE, IPSec, L2TP).

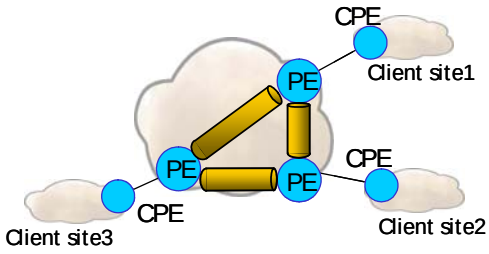

In L3 VPNs, CE and PE exchange routing information (IP reachability information). A type of L3 VPN is defined in [ROSE06]. This model is widely used by VPN SPs.

Typical L2 VPNs over MPLS are defined in [KOMP05] or [LASS05].

As the peer model is the predominant model for VPN Service Providers, AGAVE will focus on this MPLS based VPN.

14.2 Building Inter-domain VPN

The purpose of this service is to build a VPN whose sites are attached to different ASes.

There are currently four ways to create inter-domain VPNs. The first three ones -called options "a", "b" and "c"- are specified in [ROSE06]. The last one -called option "d"- is presented in [KULM06].

In each option, the AS can either be owned by the same SP (inter-AS) or by different SP (inter-SP).

14.2.1 Inter-AS VPN option "a" (back to back PE)

In this procedure, a PE router in one AS attaches directly to a PE router in another. The two PE routers will be attached by multiple sub-interfaces, at least one for each of the VPNs whose routes need to be passed from AS to AS. Each PE will treat the other as if it were a CE router. That is, each PE associates each such sub-interface with a VRF, and use eBGP to distribute unlabelled IPv4 addresses to each other.

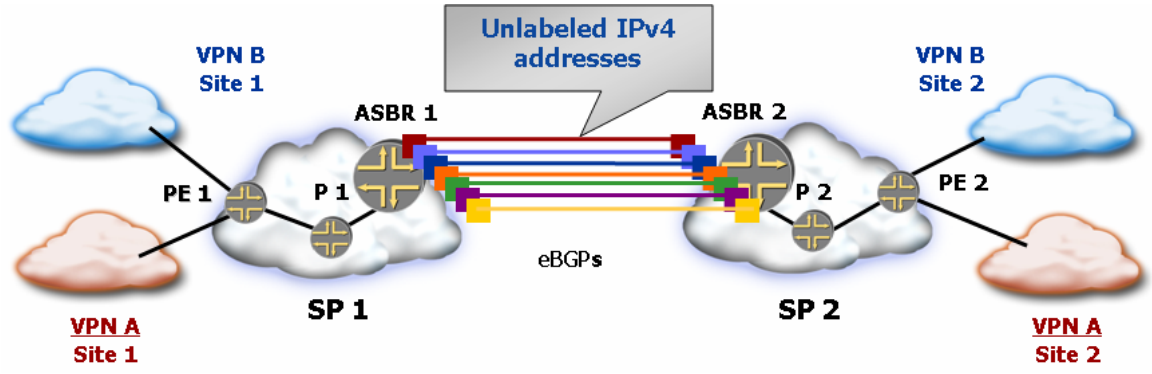

Figure 35 Inter AS VPN option a

This is a procedure that "just works", and that does not require MPLS at the border between ASes. However, it does not scale as well as the other option discussed below.

In this option, ASBRs are PE routers and therefore are VPN aware. Both PE ASBRs need to be directly connected at the IP layer i.e. by a L2 link or L2 network or possibly a L3 (IP or MPLS) tunnel.

14.2.2 Inter-AS VPN option "b" (VPN ASBR)

In this procedure, the ASBR uses Multi-Protocol Exterior BGP (MP-eBGP) to redistribute labelled VPN-IPv4 routes to an ASBR in another AS.

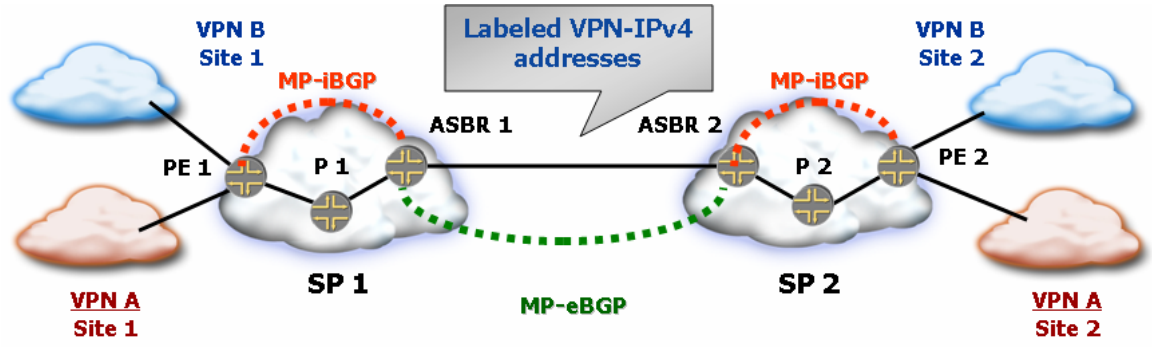

Figure 36 Inter-AS VPN option b

When using this procedure, VPN-IPv4 routes should only be accepted on MP-eBGP connections at private peering points, as part of a trusted arrangement between SPs. VPN-IPv4 routes should neither be distributed to nor accepted from the public Internet, or from any BGP peers that are not trusted. Also, there must be agreement among the set of SPs as to which border routers need to receive routes with which Route Targets.

An ASBR should never accept a labelled packet from an EBGP peer unless it has actually distributed the top label to that peer.

In this option, ASBRs are PE routers and therefore are VPN aware. Both PE ASBRs need to be directly connected at the IP layer (i.e. by a L2 link or L2 network or possibly a L3 (IP or MPLS) tunnel).

14.2.3 Inter-AS VPN option "c" (MP-eBGP multi-hop)

In this procedure, there are two eBGP peering relationships:

- a multi-hop MP-eBGP redistribution of labelled VPN-IPv4 routes between source and destination ASes.
- an eBGP redistribution of labelled IPv4 routes from AS to neighbouring AS.

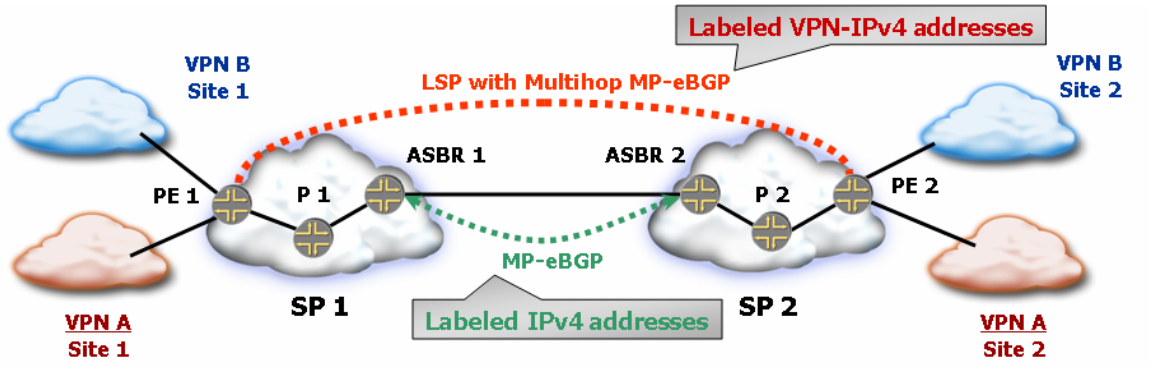

Figure 37 Inter-AS VPN option c

In this procedure, VPN-IPv4 routes are neither maintained nor distributed by the ASBRs. An ASBR must maintain labelled IPv4 /32 routes to the PE routers within its AS. It uses eBGP to distribute these routes to other ASes. ASBRs in any transit AS will also have to use eBGP to pass along the labelled /32 routes. This results in the creation of a LSP from the ingress PE router to the egress PE router. Now PE routers in different ASes can establish multi-hop MP-eBGP connections to each other, and can exchange VPN-IPv4 routes over those connections.

To improve scalability, one can have the multi-hop MP-eBGP connections existing only between a route reflector in one AS and a route reflector in another. (However, when the route reflectors distribute routes over this connection, they do not modify the BGP next hop attribute of the routes.) The actual PE routers would then only have MP-iBGP connections to the route reflectors in their own AS.

In this option, ASBRs only provide IP/MPLS transit. They are *not* PE router and are VPN *unaware*.

14.2.4 Inter-AS VPN option "d" ("a"+"b")

This option has just been proposed in the IETF and is currently an Internet Draft. This option is similar to option "a" as ASBRs are PEs with configured VRF and IP lookup performed in the VRF. It is similar to option "b" as a single MP-eBGP session is used to exchange labelled VPN-IPv4 routes and MPLS labels are used to demultiplex VPNs (MPLS labels acts as the (sub)interface of the option "a").

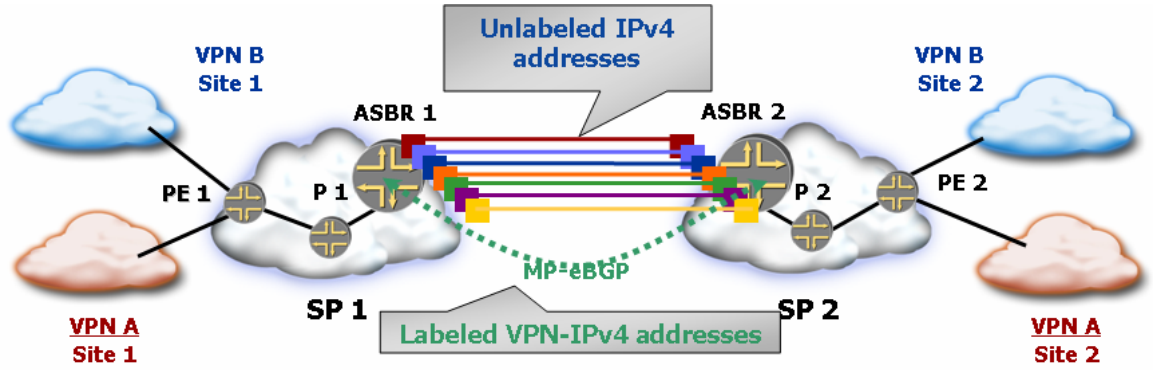

Figure 38 Inter-AS VPN option d

In this option, ASBRs are PE routers and therefore are VPN aware. Both PE ASBRs need to be directly connected at the IP layer i.e. by a L2 link or L2 network or possibly a L3 (IP or MPLS) tunnel.

14.2.5 Carrier's Carriers

Sometimes a VPN may actually be the network of an SP, with its own interconnections and routing policies. Sometimes a VPN may be the network of an SP that is offering VPN services in turn to its own customers. VPNs like these can also obtain backbone services from another SP, the "carrier's carrier", using a MPLS VPN. However, it is necessary in these cases that the CE routers and PE's VPN Routing and Forwarding tables (VRFs) support MPLS.

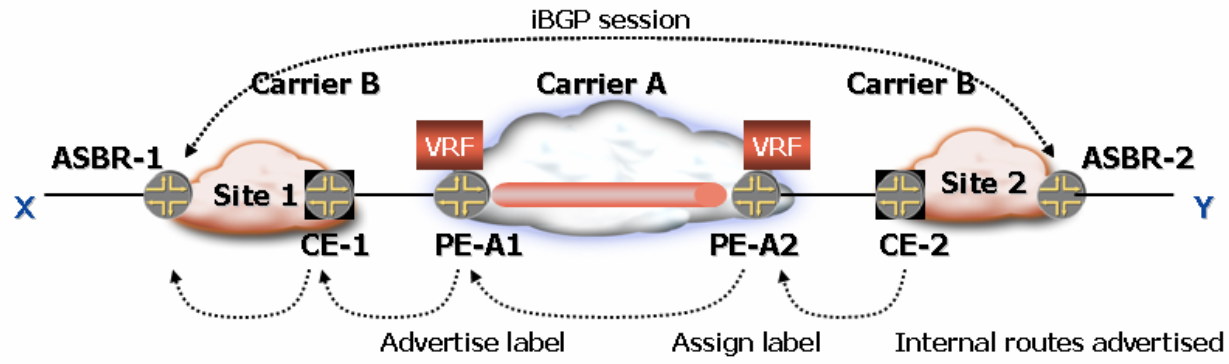

Figure 39 VPN Carrier's Carrier

In the above example, Carrier A is a VPN SP which provides IP and MPLS connectivity to its customer Carrier B through a MPLS VPN.

As Carrier B has MPLS connectivity between its sites, it can provide services over MPLS (e.g. example MPLS VPN service) to its own customers (X and Y).