AGAVE

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

IST-027609

D2.1: Initial Specification of the Connectivity Service Provisioning Interface Components

Document Identifier: AGAVE/WP2/ALGO/D2.1/public	
Deliverable Type: Report **Deliverable Nature:** External	**Contractual Date:** 30 November 2006 **Actual Date:** 14 December 2006

Editor:	Eleni Mykoniati, Algo
Authors:	*TID:* M. L. García Osma, A. J. Elizondo, J. Rodríguez Sánchez *FTR&D:* M. Boucadair, B. Decraene, B. Lemoine, J.L. Le Roux *Algo:* E. Mykoniati, P. Georgatsos *UCL.uk:* D. Griffin, J. Spencer, J. Griem *UniS:* N. Wang, M. Amin, K. H. Ho, M. Howarth, G. Pavlou *UCL.be:* B. Quoitin, O. Bonaventure
Abstract:	This document specifies an open connectivity provisioning interface to allow Service Providers to interact with underlying IP Network Providers for the provision of end-to-end IP-based services. The requirements for provisioning IP-based services in the Internet -discussed and analysed through typical service case studies such as VoIP and VPN in project deliverable D1.1- are captured in the so-called Connectivity Provisioning Agreement (CPA).
Keywords:	Quality of Service, Network Planes, Parallel Internets, Connectivity Provisioning Agreement

Copyright © AGAVE Consortium:

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

This document is the first WP2 deliverable of the AGAVE project.

The project advocates a 'clear-cut' interface between Service Providers (SP) and IP Network Providers (INP). Drawing from the requirements of typical IP-based services, such as plain IP connectivity, VoIP (Voice over IP) and VPN (Virtual Private Network), a suitable business model has been elaborated and service connectivity requirements have been captured and consolidated in project deliverable D1.1 [D1.1]. From the standpoint of an INP, a functional architecture for supporting the interactions with SPs to provide the required IP connectivity was also drawn.

Departing from the work in [D1.1] and building further on the Parallel Internets framework, this document focuses on the interface between INPs and SPs. In particular, the document specifies an open connectivity provisioning interface to allow Service Providers to interact with underlying IP Network Providers for the provision of end-to-end IP-based services.

Service Providers interact with IP Network Providers on the basis of *Connectivity Provisioning Agreements* (CPAs). A CPA allows for defining the IP connectivity requirements of a Service Provider in terms of QoS and availability guarantees in a specified scope. Further, a CPA allows for defining access control, shaping, flow forwarding and routing rules to be enforced at particular edges or across the defined scope. Through CPAs, Service Providers may also specify performance reports and notifications that wish to receive either for the assurance of their CPA or as feedback for driving their own dynamic service engineering functions.

Copyright © AGAVE Consortium, November 2006

Table of Contents

Copyright © AGAVE Consortium, November 2006

Detailed Table of Contents

Copyright © AGAVE Consortium, November 2006

List of Figures

Copyright © AGAVE Consortium, November 2006

1 INTRODUCTION

This deliverable is produced by AGAVE Work Package 2, which focuses on *Connectivity Service Provisioning*. WP2 has been setup with the following objectives:

- To specify a unified interface supporting the common provisioning and control requirements for the connectivity aspects of end-to-end IP-based services within the Parallel Internets framework, with the ultimate objective to facilitate the rapid deployment of services.
- To identify the generic networking capabilities of Network Planes and design the Network Planes management interface to support the operations of the service provisioning interface.
- To specify an overall engineering approach and select appropriate implementation technologies.
- To design and implement the components realising the operations supported by the connectivity service provisioning interface and their interactions with the underlying network through the Network Planes management interface.
- To specify test requirements for evaluating the validity of the specifications and development focusing on specific service type and business model use cases.

This deliverable addresses the first WP2 objective by specifying the external interface of the INP. The specification work documented herein, is undertaken in activity AC2.1 and is based on the results of Work Package 1 captured in deliverable [D1.1]. The *business model* identified in [D1.1] is adopted. The *requirements* and the *high-level specifications* of interactions among Service Providers (SPs) and IP Network Providers (INPs) constitute the starting point for the specification of the *external INP interface* documented in section 2 of this deliverable. Following the first WP2 objective (see above), this interface intends to capture the provisioning and control requirements for the connectivity aspects of the services studied in the corresponding *business cases*.

In particular, only the vertical interface the INP offers to the SP is studied. The horizontal interface between INPs is considered to be a subset of the vertical interface, as shown by the high level specification of the corresponding interactions (see [D1.1], sections 6.1.1.1 and 6.1.2.1). Moreover, the horizontal interface between INPs has been thoroughly studied in the past (see [MESCAL]) in the context of cascaded agreements between peer DiffServ-enabled INPs.

The specification of the external connectivity provisioning interface is work in progress. This deliverable captures the CPA specification. Future work is expected to revise and enhance current specifications, following the results of implementation and validation work. The final specification will be captured in the AGAVE deliverable D2.2, due on February 2008.

Copyright © AGAVE Consortium, November 2006

2 INTERFACE TO SERVICE PROVIDER

2.1 Introduction

Defining the IP Network Provider (INP) as an autonomous role interacting directly with Service Providers (SPs) – from network layer SPs to higher layer Application SPs (ASPs) – introduces an interface between INPs and SPs, exposing the IP connectivity capabilities of the INPs in a generic service-provisioning-aware but not service-specific way (see [D1.1]). 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.

The idea of defining an interface to clearly distinguish the operational concerns in the IP and service layer has always been thought of as a 'good practice', mainly on grounds of hierarchical system design, bringing amongst others the merits of encapsulation of lower level functions and separation of concerns. In addition to these engineering merits, the introduction of INP-SP interface advocates new business roles and therefore bears new business opportunities in the chain of service delivery in the Internet. In line with this view are emerging studies, which also 'break' the traditional role of an ISP along the lines of 'networks and services' proposing 'clear-cut' interfaces [FEAM06].

Beyond the forwarding and the QoS treatment of the traffic entering the INP's network from the SP's sites, the INP offers to the SP means to control the connectivity provisioning. Particular connectivity provisioning requirements are captured for the IP connectivity, VoIP and VPN service business cases studied in [D1.1].

The main ingredient of the INP-SP Interface is what we call the *Connectivity Provisioning Agreement* (CPA). SPs interact with INPs on the basis of such agreements. The interface provides for the necessary means to negotiate CPAs and execute the operational actions agreed in the CPA.

In a snapshot, the CPA allows for defining the connectivity requirements of the SP in terms of QoS and service availability guarantees in a specified scope. Further, the CPA allows for defining access control, shaping, flow forwarding and routing rules to be enforced at the edges or across the defined CPA scope. The SP may also specify performance reports and notifications that it wishes to receive either for the assurance of the CPA or as feedback for driving its own dynamic service engineering functions.

Our work draws from the SLS template [TEQUILA] and the provider SLS (pSLS) template [MESCAL] specifications of the IST TEQUILA and IST MESCAL projects. TEQUILA specified a service management and traffic engineering framework for *intra-domain QoS provisioning*, which prompts for standardisation of the notion of SLS, proposing a standard template to capture the IP QoS-aware connectivity services offered to end customers. MESCAL advanced the SLS template specification to incorporate *inter-domain aspects*. The provider SLS (pSLS) template models the interactions between peer providers for the provision of end-to-end IP QoS-aware connectivity services to end customers, assuming that peer providers co-operate for providing QoS guarantees in a cascaded way.

Although AGAVE CPA specification builds on the TEQUILA SLS and MESCAL pSLS, the CPA model is different from these in the following aspects: a) AGAVE focuses on the interactions between the -now distinct- IP Network Provider and Service Provider roles, while in TEQUILA and MESCAL these two roles were encompassed in the role of an ISP, b) CPA captures the connectivity provisioning requirements of SPs rather than the connectivity requirements of end customers and c) the AGAVE CPA does not depend on the strict cascaded model between peer providers as assumed by the MESCAL pSLS model.

Copyright © AGAVE Consortium, November 2006

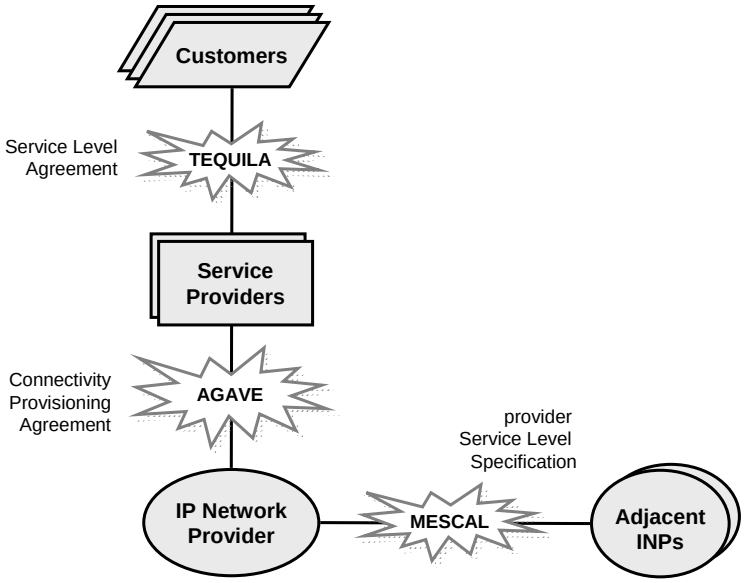

Figure 1 AGAVE CPA with respect to previous work

2.2 CPA template specifications

The CPA is specified against the information elements (clauses) shown in Figure 2[1], see details in the following sections.

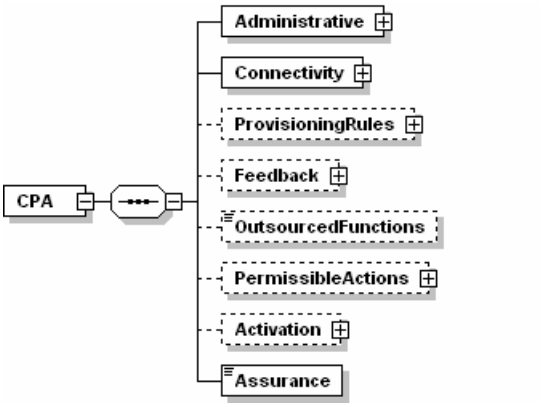

Figure 2 Connectivity Provisioning Agreement

2.2.1 Administrative

The administrative information clause may include information on the Service Provider, the formula to be used for charging, etc.

2.2.2 Connectivity

The connectivity clause captures the IP connectivity requirements of the SP. The IP connectivity is specified in terms of *connectivity modules*. A connectivity module specifies capacity and guarantees within a defined scope. The desired guarantees are modelled as *connectivity classes*. The scope is defined as connections between specific *edges*, local or remote to the INP domain. Unlimited scope is allowed and is specified as a remote edge encompassing all possible destinations.

[1] The figures are drawn with the Altova XML editor. For an explanation of the diagram model, see section 5.1.2 Content Model View (pp 135-144) of *Altova XMLSpy 2007 Enterprise Edition User & Reference Manual*, available at http://www.altova.com/documents/XMLSpyEnt.pdf.

Copyright © AGAVE Consortium, November 2006

2.2.2.1 Edges

An edge (see Figure 3) denotes the ingress or egress border link or border router where the responsibility of the INP for delivering the traffic according to the terms of the CPA begins or terminates.

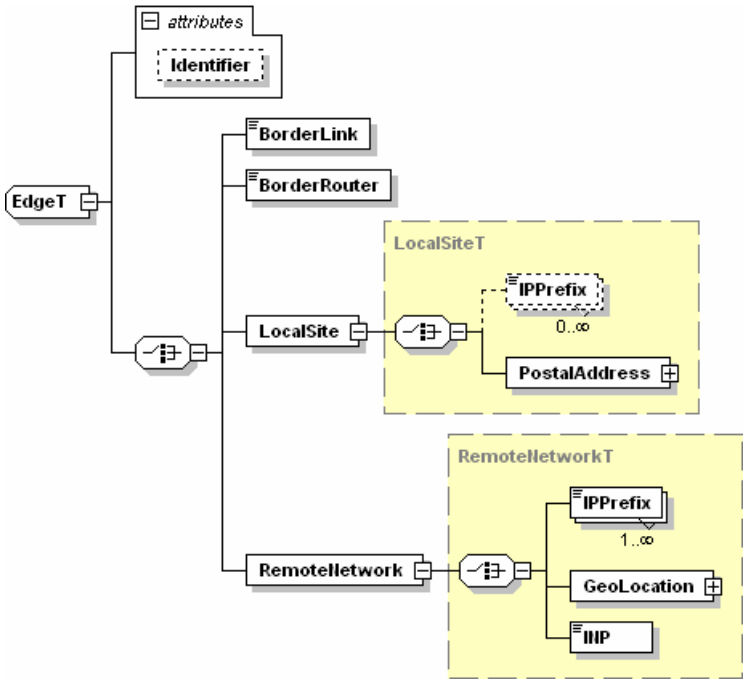

Figure 3 Edge

Edges are used to determine the scope of connectivity modules (see section 2.2.2.3). An INP may provide guarantees for reaching directly attached sites only or for remote sites too, the latter is called a multi-hop CPA (see [D1.1], section 6.1.2). An edge can thus be specified as the border link or border router either local to the INP, or of a remote INP in case of a multi-hop CPA. Instead of the INP link or router, an edge can be specified as a SP or customer site directly attached to the INP, or as a remote network represented by a set of IP address prefixes. Note that, when the CPA edge is the final destination of the SP traffic, it is transparent to the SP whether it is local or remote to the INP. However, when the CPA edge is only an intermediate node for the SP traffic, this implies that the SP has another agreement with the provider (INP or SP) connected to the remote end, hence the SP is aware of the CPA edge location.

In case of a *local site*, the INP matches the site information with some registered information to derive the corresponding border link where the site is connected. In case of multi-homed sites, a local site edge will be mapped internally by the INP to several edges (border links), as many as the connections between the site and the INP. In case of a *remote network*, the INP translates the provided geographical areas or peer INP identifiers to IP prefixes and derives the downstream NIAs and associated inter-domain links towards these remote destinations to be used to carry the SP traffic, based on the downstream NIAs and routing configuration in effect. In this case also, one remote network edge may be translated to several local border link edges, or the opposite, many remote network edges may be translated to just one local border link edge.

2.2.2.2 Connectivity Class

The connectivity class (see Figure 4) captures the guarantees on IP packet transfer performance metrics that the INP agrees to offer to the SP traffic in the context of a connectivity module.

A connectivity class includes the following attributes, corresponding to the IP packet transfer performance and IP connectivity availability metrics against which guarantees are given:

Copyright © AGAVE Consortium, November 2006

- Delay guarantees, specifying the guarantees for the one-way packet delay as measured between specific ingress and egress points crossed by the SP traffic.
- Jitter guarantees, similar to the above.
- Loss guarantees, specifying the guarantees for the packet loss probability; this is defined as the ratio of the lost packets between specific ingress and egress points and the injected packets at ingress.
- Availability guarantees, specifying the percentage of the time over an agreed measurement interval, where the above QoS guarantees are provided as agreed; availability guarantees thus capture the frequency and persistence of physical failures and/or QoS degradation caused by congestion.

It is not necessary for all above attributes to be specified. Relevant metrics have been standardised (see [RFC2679, RFC2680]). However, the metrics supported by each INP may vary depending on its policies and capabilities.

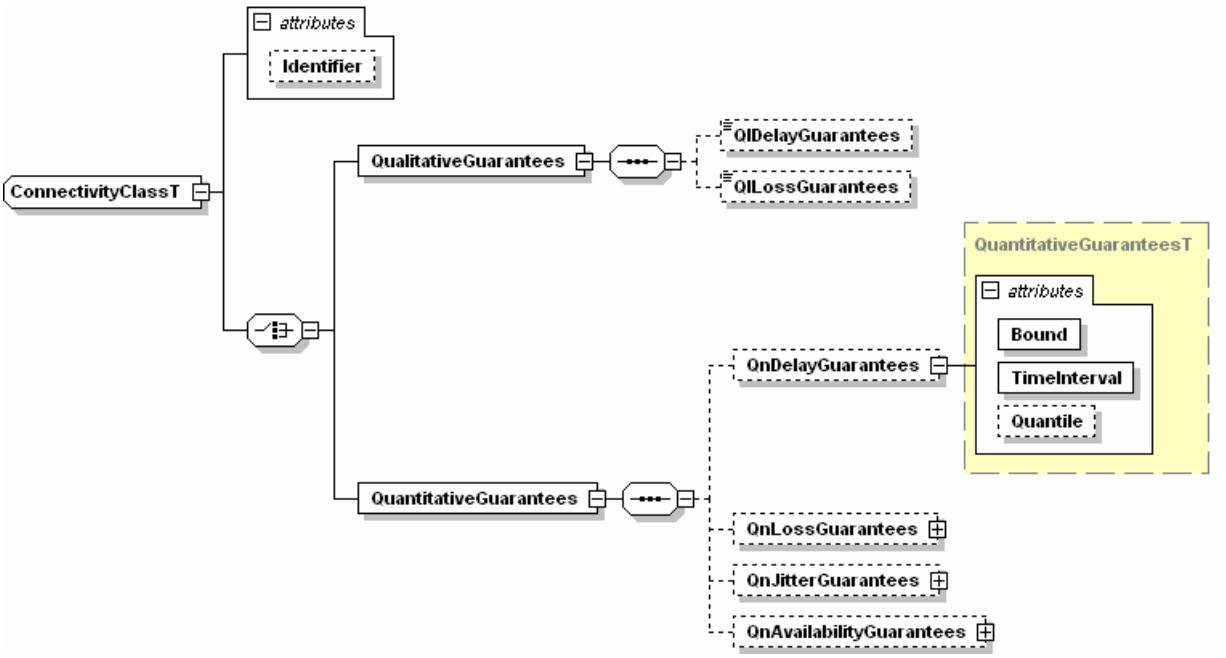

Figure 4 Connectivity class

The following aspects underlying the semantics of the above attributes are worth noting:

The following types of guarantees are distinguished: quantitative and qualitative. The guarantees to a particular metric are said to be quantitative, if they can be expressed in quantitative, numerical, values. Otherwise, they are said to be qualitative; possible qualitative values, as appropriate as per performance parameter, may include: high, medium, low or red, yellow, green. The quantification of the relative difference between the qualitative values is a matter of provider's policy e.g. 'high' could be twice good as 'medium', which in turn is twice as good as 'low'.

Quantitative performance guarantees are expressed as maximum (worst-case) bounds or as (sets of) percentiles or inverse percentiles, indicating also the granularity period of the associated measurements. The meaning of the values of qualitative performance guarantees and/or their relative difference should be clear to the SPs, while it should be backed-up with relevant historical performance data.

Each connectivity class specified in a CPA will be mapped to appropriate Network Planes and Parallel Internets internally in the INP.

Copyright © AGAVE Consortium, November 2006

2.2.2.3 Connectivity Module

The connectivity requirements of the SP are specified in terms of connectivity modules. The connectivity module associates a connectivity class (guarantees) to a specified scope and capacity (see Figure 5).

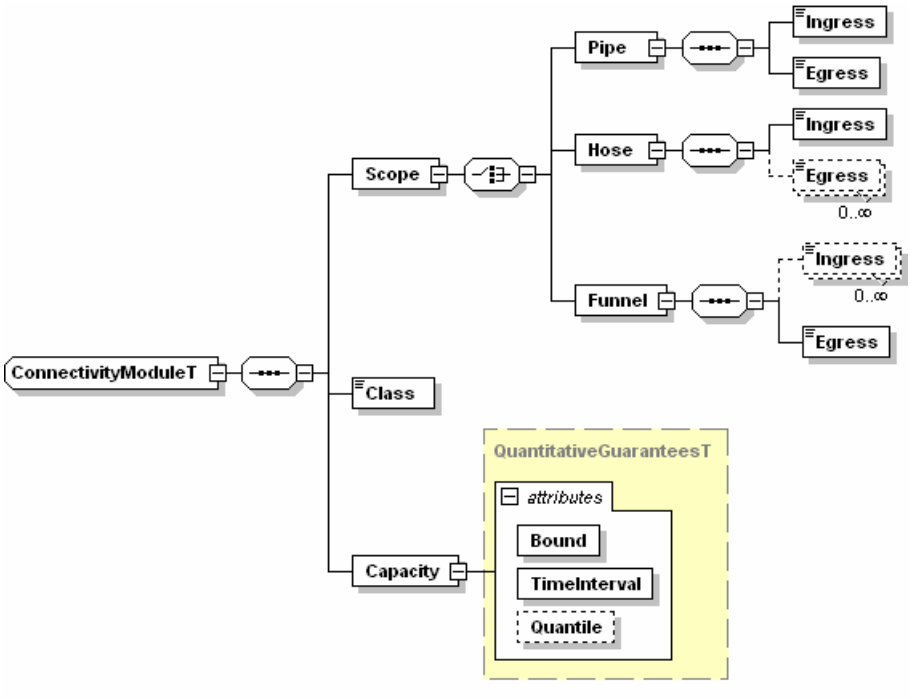

Figure 5 Connectivity module

2.2.2.3.1 Scope

Scope explicitly identifies the geographical/topological region over which liability for the connectivity class guarantees ends, by indicating the boundaries of that region in terms of edges. It includes the following attributes:

- Ingress edge, indicating the entry point of the region over which the connectivity module is to hold
- Egress edge, indicating the exit point of the region over which the connectivity module if to hold

The following combinations of Ingress, Egress values are allowed:

- (1,1) implying an one-to-one communication; we call the connectivity module a pipe
- (1,N) one-to-many communication (N>1); we call the connectivity module a hose
- (1,any) one-to-any communication; we call the connectivity module an unspecified hose
- (N,1) many-to-one communication (N>1); we call the connectivity module a funnel
- (any,1) any-to-one communication; we call the connectivity module an unspecified funnel

Because connectivity modules are assumed unidirectional, the above taxonomy excludes the many-to-many communication (M, N); either ingress or egress attributes must be specified to exactly one interface identifier. Many-to-many communication can be achieved at the level of CPA, where a number of connectivity modules are combined.

In case where the specified edges are remote, their mapping to the provider's domain boundary links is subject to the NIAs and routing decisions in place. Hence, internally in the provider and transparently to the agreed CPA, traffic to remote sites may be merged over one boundary link or split to many boundary links turning a hose to a pipe, or a pipe to a hose, etc.

Copyright © AGAVE Consortium, November 2006

Usually, one of these attributes corresponds to an interconnection link of an SP site to the reference INP, while the other attribute is left unspecified or set to a set of destinations (remote network), the interconnection link of another SP site to another INP, or the boundaries of another INP. As an example, in the case of an Internet SP the value of the ingress would be the SP's interconnection points and the value of the egress would be left unspecified denoting "any"; the latter could be refined to denote the interface of a particular inter-domain link by the traffic engineering functions, however, this is an internal matter, being not subject of agreement. In the case of a VPN SP, both ingress and egress would be clearly specified.

2.2.2.3.2 Capacity

The capacity clause determines the SP traffic volume that can be supported at the guarantees of the specified connectivity class in the specified scope.

When the scope of the connectivity module is a pipe (see section 2.2.2.3.1) the INP must verify that the required resources are available from the ingress to the egress. When the scope is a hose, the INP must verify that the required resources are available from the ingress to any of the egress points. INP is forced to allocate resources equal to the overall capacity across all ingress-egress pairs, resulting in a significant resource underutilisation. This can be smoothed out, by constructing tree paths from the ingress to the egresses.

2.2.3 Provisioning Rules

Provisioning rules include rules on access to the specified connectivity, on the forwarding of flows across the edges of the CPA, routing rules for constructing the paths between the CPA edges, and finally rules for shaping the traffic at the egresses. Other types of rules may be included in the future.

2.2.3.1 Access Rules

Access rules (see Figure 5) specify how data flows are treated within the CPA, including policing at the CPA ingresses, assignment to the connectivity class for enjoying the associated guarantees, and marking policies at the CPA egresses. An access rule is defined for a particular micro-/macro-flow entering from one or a set of the CPA ingress points.

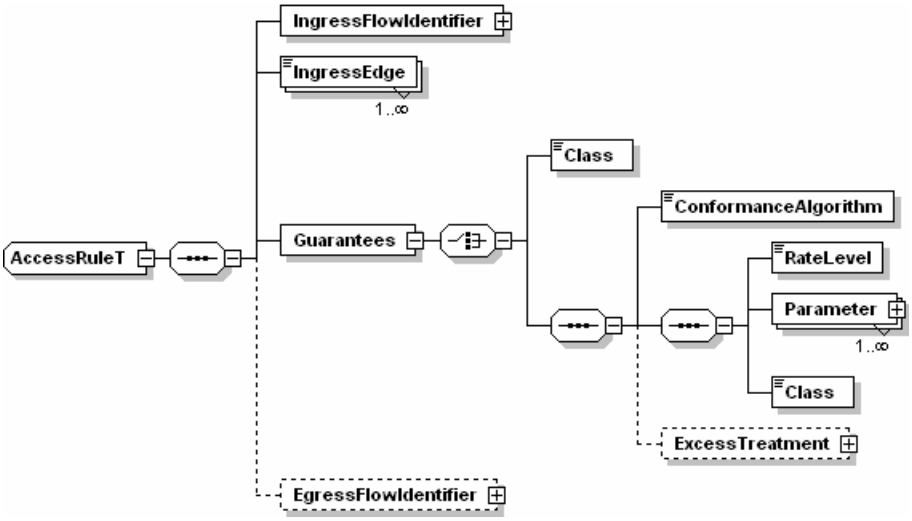

Figure 6 Access rule

2.2.3.1.1 Ingress flow identifier

The ingress flow identifier sets the classification rules identifying the stream of IP datagrams constituting the flow to which the access rule is to apply. Classification is performed based on the IP

Copyright © AGAVE Consortium, November 2006

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.

2.2.3.1.2 Guarantees

The access rule determines the guarantees the identified flow is entitled to and the restrictions the flow must adhere to for getting the guarantees. Note that the access rule merely controls how much volume of which flows will gain access to the capacity associated to the connectivity modules; the access rule does not per se implies that all packets admitted following the defined access rules will get the associated guarantees, not if the rate of the overall injected traffic exceeds the capacity of the corresponding connectivity modules (see section 2.2.2.3).

The identified flow can be associated with a connectivity class (see section 2.2.2.2) unconditionally for all received packets, or subject to a conformance traffic profile. In the latter case, the following attributes are specified:

- Traffic conformance algorithm, specifying the mechanism used to unambiguously identify the packets complying with the traffic conformance criteria and those which do not, called the "in" and "out" of profile packets, respectively. Examples of traffic conformance algorithms are: leaky bucket, token bucket, combined token bucket with peak, a two-rate three-colour marker scheme [RFC2698] and an MTU-based scheme.
- Traffic conformance algorithms may allow for setting multiple levels of traffic rate conformance. Each traffic conformance level is characterised by associated conformance criteria in terms of rate (bandwidth) thresholds, captured in traffic conformance parameters like peak rate, token bucket rate, bucket depth and maximum transfer unit (MTU).
- The traffic that conforms to a particular rate level is assigned to a connectivity class.
- The treatment of the traffic in excess of the highest rate level is specified. Excess treatment may be dropping (default), shaping, or gaining access to a qualitative connectivity class. Note that the rate of excess traffic is unlimited, hence it is senseless to assign it to a connectivity class designed to deliver quantitative guarantees.

2.2.3.1.3 Ingress edges

Each access rule applies to one or more ingress CPA edges. If the CPA edge is a border link then this implies that policing rules corresponding to the traffic conformance clause will be configured at the corresponding border interface. If the ingress CPA edge is a border router, then the same policing rules will be configured at every external input interface. As a result, if a flow is distributed to more than one path from the upstream domain, it would gain access to a rate higher than the highest conformance level, as many times as the number of CPA edge border interfaces it is mapped to.

At every border router associated with an ingress CPA edge, forwarding will be configured so as to forward the packets of the packets conformant at each level to the Network Plane and Parallel Internet corresponding to the connectivity class of this conformance rate level. A validity check on the CPA specification will require that the ingress edge and connectivity class associated with a flow have a match to a specified connectivity module.

2.2.3.1.4 Egress flow identifier

Egress flow identifier captures the requirements for the marking and/or tunnelling identifiers to be applied to a certain flow at the CPA egresses. VPN traffic for example is expected to require DSCP transparency (see [D1.1], section 5.3.1.1), which forces the INP to maintain the ToS/DSCP values as in the original packets received at the CPA ingresses. Other flows may require re-marking with a particular ToS/DSCP value, because so it is expected by the other end following the CPA egress. Note that, delegating remarking at the CPA egresses raises interoperability and scalability issues for multi-hop CPAs (see [D1.1], section 6.1.1.5.1). By default when no egress flow identifier is specified, all INPs in the path between the CPA ingress and egress are free to remark SP's traffic.

Copyright © AGAVE Consortium, November 2006

2.2.3.2 Forwarding Rules

The clause of forwarding rules includes per flow route selection rules, specifying the egress CPA edge where the defined flow should be directed to. Such rules allow for overriding the routing of the INP, enabling the SP to implement its own routing logic over a logical topology where CPA egress edges are not the final destinations but intermediate nodes to paths controlled by the SP. Forwarding rules could be specified by overlay SPs, or by VPN SPs, etc.

2.2.3.3 Routing Rules

Routing rules determine constraints and preferences for constructing the path between the CPA edges, e.g. exclude a particular AS from the inter-domain path. A routing rule differs from a forwarding rule in that a forwarding rule indicates the egress CPA edge for a flow at a particular ingress edge, while a routing rule determines how the logical link between the CPA ingress and egress edges must be constructed.

2.2.3.4 Shaping Rules

Shaping rules determine profiles for shaping the SP traffic at a particular egress CPA edge. When the CPA edge is a remote location outside the domain of the INP, the INP must delegate its enforcement to the final INP in the downstream path where the egress CPA edge belongs. Identification delegation issues may arise in this case (see [D1.1], section 6.1.1.5.1).

2.2.4 Feedback

To perform fault and performance management each SP requires feedback from the network, especially in the case of qualitative guarantees. Feedback requirements are captured in the following clauses (see Figure 7):

- Monitoring tasks are specified in terms of the metrics (see [D3.1], section 4.2 for details) and data collection attributes like granularity, sampling frequency etc. The scope of a monitoring task may be limited to the INP domain or it may include inter-domain statistics extending as far as the final destinations or the CPA egress.
- Notifications and reports, determine the monitoring reports and/or alarms that the INP must produce and deliver to the SP periodically or on-demand. The notifications and reports are specified in terms of the metrics defined in the monitoring tasks clause, setting reporting frequency, alarm thresholds, etc. Applicable reports/alarms are related to network performance, failure incidents, security attacks, troubleshooting logs etc.

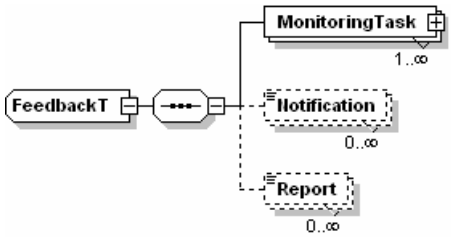

Figure 7 Feedback

2.2.5 Outsourced Functions

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.

Copyright © AGAVE Consortium, November 2006

2.2.6 Permissible Actions

The permissible actions clause (see Figure 8) determines the actions that the SP is entitled to invoke with respect to the particular CPA, for adding, deleting or modifying connectivity entries, provisioning rules or feedback requirements.

Each action is associated to one availability and one invocation profile. The action availability profile captures the guarantees and the conditions for authorising a particular action, e.g. the time of the day where this request can be performed, the probability for a request to be granted, restrictions on how frequently the action can be performed, restrictions on time to respond for the INP etc. The invocation profile captures information on how the action is to be communicated between the SP and the INP, such as the invocation protocol, and who is entitled to perform the action, allowing the SP to specify user authentication for securing and restricting the access to the CPA permissible actions.

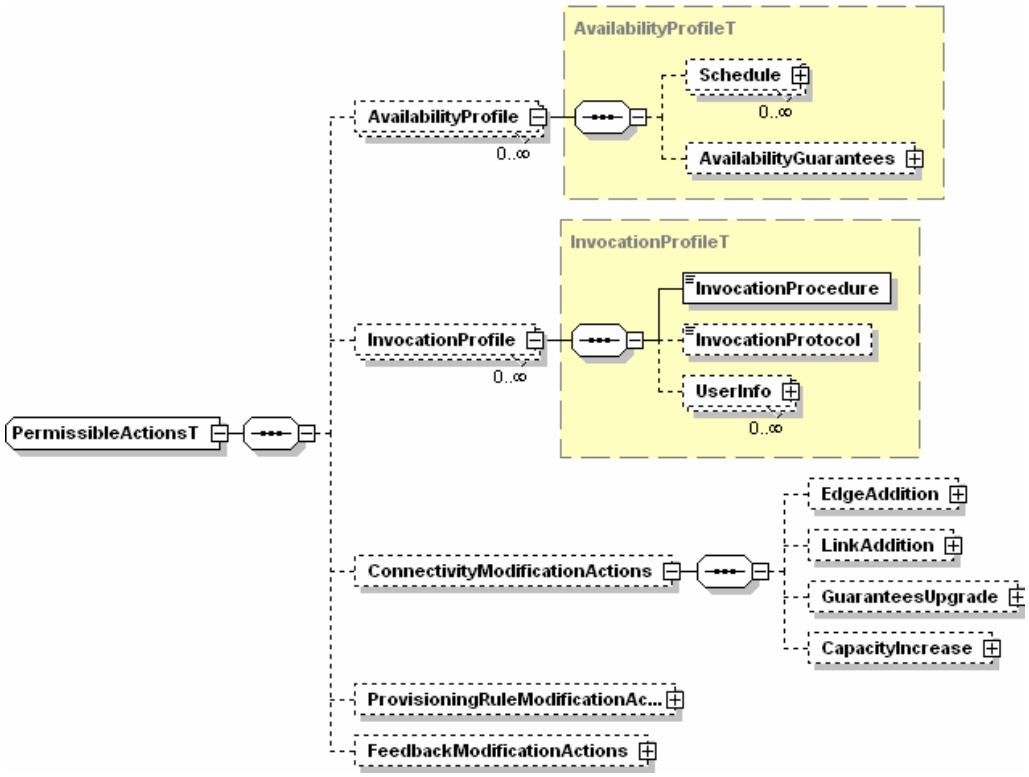

Figure 8 Permissible actions

Typical examples of connectivity modification actions are the expansion of the CPA scope by adding new local or remote edges, the increase of the CPA connectivity density by adding new links between existing edges, the upgrade of guarantees associated to existing links and the increase of capacity.

2.2.7 Activation Info

Activation information specifies potential interactions that need to take place between the INP and the SP, beyond internal INP configuration, to complete the CPA activation. 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 CPA edges, configuring authorisation to invoking probing for CPA assurance, etc.

2.2.8 Assurance

The assurance information clause determines the verification methodology and sets the *Key Performance Indicators* (KPIs) and other applicable parameters for assessing the conformance of the SP and the INP to the agreed terms and conditions. Similarly to the specification of the feedback requirements (see section 2.2.4 above), notifications and reports may be associated to each KPI to

Copyright © AGAVE Consortium, November 2006

notify the SP for service degradation. Penalties may be associated to service degradation scenarios in terms of the specified KPIs.

Note that in general, assurance monitoring and reporting is not identical with the required feedback. Assurance specifies the requirements for assessing the conformance to the CPA, while feedback may involve providing additional information and in different time scales, for either a compliant or a non-compliant CPA.

As both the SP and the INP are expected to perform monitoring for verification and assurance of the CPA, for the results to be comparable, clock synchronisation may be required (see [D3.1], section 3.4). The SP also specifies the probing facilities to which it requires access for performing its own verification measurements in terms of protocols and probing scope, within the scope of the specified connectivity modules.

Copyright © AGAVE Consortium, November 2006

3 REFERENCES

[D1.1] Boucadair M. et al., *Parallel Internets Framework*, AGAVE Deliverable D1.1, September 2006.

[D3.1] Wang N. et al., *Initial Specification of Mechanisms, Algorithms and Protocols for Engineering the Parallel Internets and Implementation Plan*, AGAVE Deliverable D3.1, December 2006.

[FEAM06] Feamster N., Gao L., Rexford J., *How to lease the Internet in your spare time*, Georgia Tech Technical Report GT-CSS-06-10, http://www-static.cc.gatech.edu/~feamster/papers/cabo-tr.pdf, August 2006.

[MESCAL] Wang N. et al., *Final Specification of Protocols and Algorithms for Inter-domain SLS Management and Traffic Engineering for QoS-based IP Service Delivery*, http://www.ist-mescal.org/deliverables/d1.3_finalv2.pdf, MESCAL Deliverable D1.3, June 2005.

[RFC2679] Almes G. et al., *A One-way Delay Metric for IPPM*, IETF RFC 2679, September 1999.

[RFC2680] Almes G. et al., *A One-way Packet Loss Metric for IPPM*, IETF RFC 2680, September 1999.

[RFC2698] Heinanen J., Guerin R., *A Two Rate Three Color Marker*, IETF RFC 2698, September 1999.

[TEQUILA] Damilatis T. et al., *Final Architecture, Protocol and Algorithm Specification*, http://www.ist-tequila.org/deliverables/D3-4b.pdf, TEQUILA Deliverable D3.4, part B, October 2002.

Copyright © AGAVE Consortium, November 2006