Market Managed Multi-service Internet

M3I

European Fifth Framework Project IST-1999-11429

Deliverable 1
Requirements Specification
Reference Model

The M3I Consortium

Hewlett-Packard Ltd, Bristol UK (Coordinator)
Athens University of Economics and Business, GR
BT Research, Ipswich GB
Eidgenössische Technische Hochschule, Zürich CH
Forschungszentrum Telekommunikation Wien, A
Technische Universität Darmstadt DE
Telenor, Oslo NO

For more information on this document or the M3I Project, please contact:

Hewlett-Packard Ltd,
European Projects Office,
Filton Road,
Stoke Gifford,
BRISTOL BS34 8QZ,
UK
Phone: (+44) 117-312-8631
Fax: (+44) 117-312-9285
E-mail: sandy_johnstone@hp.com

Document Control

Title: Requirements specifications; Reference model
Type: Public Deliverable
Author: Ragnar Andreassen / Telenor R&D (editor)
e-mail: ragnar-oivind.andreassen@telenor.com

Origin: Telenor R&D
Doc ID: \deliverables\WP2\m3idel01v7_1.doc

AMENDMENT HISTORY

Version	Date	Author	Description/Comments
V 7 Final	06-July-2000	M3I	Editorial updates
V 7.1	04-Feb-2002	M3I	Released for Public distribution

TABLE OF CONTENTS

1 Introduction

In the later years, and especially since the work of Varian, MacKie-Mason and others on Internet pricing (see e.g. the collection in [13]) there has been a general interest in the notion of dynamic pricing for network services. The term is probably a misnomer, as all prices that can be observed are dynamic in the sense that they vary over time. The essence lies however in the time scale of variation; in dynamic pricing the typical time-scale of variation is reduced by orders of magnitudes, from years and months to seconds or even milliseconds. The rationale is to improve economic efficiency of the network by realising the economic equilibria where demands and supplies are balanced through the adjustment of prices. For the establishment of such equilibria in a network, variations in prices will have to match variations in demand.

In order to best achieve the above praiseworthy goal, a market mechanism for network services should exist. The project M3I (Market Managed Multiservice Internet) is aimed at investigating how such markets can be organised, and the implications over a broad range of having them. Given that dynamic pricing gives more economically efficient networks, there are still questions to be asked with regard to implementation complexity, user acceptance and relations between stakeholders in a network-marketplace.

The participants of the M3I-project expect that market managed networks will not necessarily complicate service provision relative to other means of achieving market- and service differentiation; on the contrary there exist potentials for significant simplifications. The claim of the existence of such a double gain remains however yet to be validated. This report represents a first step for the project in that direction.

The report concerns requirements. Requirements need qualifications: who pose them to whom with regard to what. The emphasis of the project is on charging for services, so the everyday notions of “user”, “provider” and “system” are used to introduce a proposed set of essential requirements applying in general to service provisioning in that context. In the Internet, services to end-users are however as a rule established at the combined effort of a set of cooperating stakeholders. To make the results applicable and plausible, it will be necessary to understand important implications of stating general requirements to the relations between the stakeholders that together provide the services.

A role model is developed and applied to achieve the above. The concept of a role is defined on the basis of which are considered as the important functions and interactions in a market-managed multi-actor environment. In the model, stakeholders as they can be observed are then decomposed into a set of roles that are classified according to their location in the overall value chain.

The final result of the report is a set of specific scenarios describing stakeholders and their relations, using the above model. The scenarios are designed such that the relevant general requirements that were posed initially can be met.

The structure of the report follows broadly the above presentation. In chapter 2 is presented the project’s proposition of general requirements. The role model is presented and expanded in chapter 3, including a walk-through of several current European and US Internet-service stakeholders. It is believed that in an environment of dynamic pricing and cooperating stakeholders, certain general functions will be useful. We call these clearinghouse, dynamic price handler and risk broker. Their general properties and various implications of introducing the functions are explained in chapter 4. Our specific scenarios are given in chapter 5. These will form a basis for several further activities within the project. Finally, it is not enough just to present the scenarios; it is also necessary to relate them to the general requirements posed in the beginning of the document. This is the purpose of the last chapter 6.

Before proceeding with the substance of the presentations, a note on scope is relevant. The M3I project is concerned with providing differentiated and differentiating services through market forces in multi-actor environments. This includes the mechanisms in the network or the end systems required for carrying out the service differentiation, the organisation of charging for such services, the mechanisms needed to realise efficient market feedback control loops, and interaction mechanisms between cooperating parties in terms of service provisioning, charging and clearing.

The overall term "charging" is here (compliant with definition in [14]) used as a summary word for the overall process of metering resources, accounting their details, setting appropriate prices, calculating charges, and providing a fine-grained set of details required for billing.

The customer centred process of billing, defined as collecting charging records, summarising their charging content, and delivering a bill or invoice, including an optional list of detailed charges to users (definition also from [14]), is specifically not within the scope of this project. Neither are concepts and systems for payment and E-commerce considered as a part of the project scope. We are however aware that work within the project scope depend on, and may have consequences for systems dealing with billing and payment. To the extent the project or its follow-up activities performs realistic user trials, a rudimentary billing function will be needed. Also, aspects of billing may have to be considered in order to specify clearing functions as described later in this document. A specific discussion on demands payment systems pose on systems within the project scope is added in an appendix.

2 General requirements

As stated in the introduction, the common and informal concepts of “user”, “provider” and “system” will be used when expressing the M3I project’s proposition of a set of essential requirements. Further expansions on stakeholders and their interrelation will be given in the next chapter.

The general requirements are divided in “motivating requirements” and “extended requirements”. This divisioning is not really motivated through the thematic content the various requirements, but follows from the process leading up to production of this report. The motivating requirements are the ones that brought the M3I project together in the first place, and it has been found convenient keep them together in order to track their fate. In addition to these, several points not discussed in the early process, but nevertheless considered essential, are added in the section on extended requirements.

2.1 Motivating requirements

The overall goal of project M3I, as stated in its “Description of Work” is to enable ISPs to explore sophisticated charging options and business models with their customers.

We believe that end-users will value the following improvements in the internet service they receive:

U1: the ability to instantaneously increase quality of service (QoS) by accepting different charging rates;

U2: more effective competition in a differentiated services market;

U3: real-time feedback and validation of charges.

We believe that ISPs will be interested in providing the following improvements to the Internet service they provide:

P1: the ability to change tariffs and easily communicate them to the end users within seconds;

P2: the ability to hold current QoS in the presence of bad congestion effects by communicating price changes in real-time to customers;

P3: the ability to charge differentially for applications requiring differing QoS levels, or multicast.

The overall motivation for our work is thus to provide the ability to investigate the acceptability of the proposed improvements for user, and feasibility of the proposed improvements for the provider, specifically through the ability to test whether and to what extent:

D1: the demand for Internet services, including various QoS levels can be managed effectively through a pricing mechanism;

D2: customers can flexibly access both high and low quality services, depending on their particular application needs, instead of being limited to a single best-effort service as in the current Internet;

D3: end users in corporate organisations can exercise similar choice, but constrained by the policy of the party that is paying;

D4: ISPs can recover the costs of new services, such as voice and video, that are currently provided by different infrastructures, and hence increase social efficiency by exploiting economies of multiplexing and scale, which in turn will also provide for increased network revenue;

D5: simple and scalable extensions to current technology can provide the correct incentives for the economically efficient and uncongested operation of the Internet.

2.2 Extended requirements

In addition to the above motivating requirements, the following requirements of a general kind are considered essential. It should however be noted that they are considered essential relative to current knowledge and conventions. User requirements in particular might evolve and relax over time.

2.2.1 User requirements

User requirements are applicable in the general context of service deployment. Of particular interest are aspects related to relations between services and charges in situations where prices and service properties may vary dynamically, and where a service may be a composition of several chargeable service components.

- Predictability of Charges
  Users may want to be able to predict the costs of using a particular communication service. Therefore, an exact a-priori specification of communication charges would be desirable. However, if this cannot be fulfilled, a set of weaker demands can be sufficient, for example the ability to know or configure a maximum price or spending cap. Finally, it must be prohibited that a user is charged a higher price than previously announced, without giving his explicit approval.

- Transparency and Accuracy of Charging
  To find out how much is spent for which application and which factors contributed to the final charge, users need the ability to determine the charges of a particular application flow. This information is very important for a user in order to decide whether a certain service and its quality offer good value for the price. Furthermore, it is a requirement that in principle, an auditing mechanism exists, allowing an authorised third party, for example a law enforcement or consumer protection agency to investigate details of service provision.

- Convenience
  Charging components should not make the usage of communication services much more difficult. All charging mechanisms themselves as well as the final bill based on the information gathered by the charging system must be convenient for its users. Hence, it must be possible for users to define “standard charging behaviour” for their applications so that they are not bothered with details during the start up of an often-used application. On the other hand, they should be able to change such a description easily to have control over their expenditures.

- Applicability
  Charging for services should not require modification of users’ applications.

- Reliability of Invoices and Payment Methods
  Furthermore, most users want to have as few separate bills as possible, i.e., have contracts and according business procedures with only one provider. These bills must incorporate all details as agreed upon in advance and must be reliable.

2.2.2 Provider Requirements

Provider requirements encompass general features and characteristics that are related to the provision of a service and the provider's placement in an open market environment. Therefore, besides the views on the technical feasibility aspects the support of various business models forms the essential prerequisite on a market-driven approach.

- Technical Feasibility
  The charging approach and its mechanisms must be implementable and operable with low effort and low cost. Specifically, the added overhead for communication due to additional information transmitted between senders, network nodes, and receivers, and also for processing and storage purposes especially in network nodes, e.g., to keep and manipulate charging information, must be as low as possible [2].
- Variety of Business Models.
  The business of providing network service over packet-switched networks must be sustainable and profitable to attract the necessary investments into the infrastructure. It is unreasonable to expect all service providers to adopt exactly the same business model and strategies. Therefore, charging mechanisms must be flexible enough to support a large variety of business models and interoperate between multiple network domains employing different models. A charging system must also be flexible enough to handle different pricing strategies, for example during peak and off-peak times.
- Interconnect scenarios.
  M3I solutions should be applicable to a wide range of interconnect scenarios. At least two models have been proposed in recent litterature [1], [4], [5]: One, in which each intermediate ISPs expects one of the far end-customers to pay its local charges, and one in which intermediate ISPs only charge their neighbours (the status quo?). Both options should be covered by the M3I architecture, although main focus will be on models with settlements between neighbours (edge pricing).
- Bundling of ISP and application charges
  Should be facilitated, at least for some applications. Both the cases where application charge is bundled with communication charge by a connectivity provider, and communication charge is bundled with application charge by an information provider should be catered for.

2.2.3 System Requirements

The set of system requirements is related to the design of the M3I systems for Charging & Accounting and price setting: They form the basis for the evaluation of the implementation of these systems with respect to their efficiency. At the same time, they also address the important issues of adapting the systems designed to the "moving targets" in economic networking research. For instance, the variety of application scenarios may change and will require different models of charging to be incorporated.

Finally, the area of protection against fraud and intrusion is essential in an open market. The legal security aspects are studied with less emphasis within M3I.

- Flexibility
  When information is transmitted from a sender to one or several receivers, the flow of value associated with this information can be (1) in the same direction as that of the data flow, (2) in the opposite direction, or (3) a mixture of both because both sides benefit from the information exchange. For example, in the first case, the sender transmits a product advertisement, in the second case, the receiver retrieves a movie for playback, and in the third case both sides hold a project meeting via a videoconference system. To support these different scenarios, a charging architecture must provide flexible mechanisms to allow the participants in a communication session to specify their willingness to pay for the charges in a variety of manners. Senders must be able to state that they accept to pay for some percentage of the overall communication costs or up to a specified total amount. Similarly, receivers may state what amount of costs they will cover. Additionally, charging mechanisms must allow to flexibly distributing communication charges among members of a multicast group. A number of cost allocation strategies can be found in [4].

- Fraud Protection and Legal Security
  One of the most important requirements set by users of the systems is protection against fraud, i.e., that end users do not have to pay for costs they have not incurred and that no one can misuse the system. The fears of users are that a provider may cheat or that other users may use their identity or derogate from them in any other way. Providers want to be sure that users indeed pay for the used service. A prerequisite against fraud is technical security, such that users cannot damage, misuse or intrude the provider's communication systems. Finally, legal security denotes the demand that in case of a failure, there is enough information to determine responsibility for it.

- Real-time transactions
  Certain payment methods pose demands on the timely execution of charging transactions. The specific consequences on M3I-systems in realising pre-paid and pay-as-you-go methods of payment are discussed in the appendix of chapter 9.

- Openness
  When evaluating engineering proposals, the openness of each will be one of the primary factors to decide between them. It is believed that this is a commercial imperative of each of the industrial partners in the consortium, and is generally understood by all partners to be required by the Internet industry at large.

  Openness, in M3I's context, covers two related factors: build-time openness and run-time openness. Build-time openness means that at least the technical interfaces to the system are published and possibly standardised. It may mean an implementation is also published. Run-time openness means that the running system is designed to allow other commercial entities to operate services that use the system in order to operate their services.

  A market-managed system intrinsically exposes a charge to cover the cost of its operation. However, this doesn't automatically mean it is highly run-time open. It may only provide a useful service to 'retail' users. To score well on run-time openness, the system must be designed so that other services can use it on a 'wholesale' basis in order to operate their own services. Further, it should be designed to ease commercial innovation by those service providers using it.

Run-time openness' is sometimes confused with 'no access control'. For clarification, this is incorrect. An open system can still be designed to be able to deny access. It is also possible for a run-time open system to be operated closed as a matter of policy. It is the design that is judged to be open, not the policy.

Finally, openness is a qualitative measure of a system. It seems difficult to measure quantitatively, and it is also not a discrete yes/no quality.

2.3 Tensions between requirements

In the long list of requirements of the preceding sections, certain tensions have been identified. In the sequel some of these tensions will be further discussed.

2.3.1 Service quality versus charge

One of the most fundamental tensions is the one existing between the quality of a service in a broad sense, and the charge for a service. This tension is particularly present in the section above on user requirements, where requirements such as predictability and transparency will ultimately be associated with a cost. It *is* essential that such user requirements can be provided for, on the other hand it should be understood that such qualities come at a price. Users who are not interested in paying that price should have the option of relaxing their service requirements.

2.3.2 Dynamic charging versus convenience

The M3I partners want to investigate dynamic charging for the purpose of economically optimised load management. This goal is possibly reached only through frequent price updates followed by frequent customer reactions. This conflicts with the user requirement of convenience and simplicity. An end-user cannot be expected to react too often to fluctuating charges and quality. In order to resolve the conflict between end-users' restricted willingness or ability to relate directly to dynamic charging and requirements on dynamic charging for load management and economic efficiency, we propose the development of a "dynamic price handler". This will, in the ISP's eyes, replace the customer by a fast-reacting agent which is controlled by the customer through a more high level strategy.

2.3.3 Dynamic charging versus predictability

In a dynamic pricing environment, charges for services are dependent on unpredictable fluctuations in demand. There is therefore a fundamental tension between the user requirement for predictability of charges and the requirement for economic efficiency through dynamic pricing. A relaxation of this tension is possible by introducing the concept of risk brokering. A risk broker is a party that offers fixed-price service classes to the user at its own risk, using the dynamically priced underlying communication services. The risk broker's service classes may be refinements of the underlying ISP's service classes. The concepts of risk brokers and dynamic price handlers will be further discussed later in the document.

2.3.4 Diversity of service provisioning policies versus interoperability

Any one market-controlled networking solution must work within a network like the Internet, where management is highly federated. Over-provisioning, trust or coercion could be used instead of charging to allocate detailed resources. Even if market-control were used, we require that the commercial strategies of each provider should be independent from any need for standardisation such that commercial innovation should not be constrained in any way. A good market-control architecture must inter-work with networks adopting these other approaches.

Given everyone has to effectively relate to the discipline of the market, this problem reduces to one of time-scale. We are investigating dynamic pricing approaches, whereas connected networks might only see the market discipline at a coarser granularity in time. The general answer to such a problem is that any network operating a fine-grained approach will simply refuse to interconnect with one at a coarser granularity unless it can charge a premium on every longer-term charge. This would take account of the need to over-provide to avoid the otherwise increased risk of control failure. This would cover the cost of over-capacity or other coarse-grained controls such as diffserv shapers on the traffic from the less-controlled network. We must also demonstrate how it is possible to operate fine-grained market controlled mechanisms between 'market-enabled' networks separated by transit networks that have more coarse-grained control.

With regard to the 'lesser' problem of allowing innovation between networks that are operating market controls, the M3I architecture chooses to use edge-pricing as its main foundation, which helps to isolate business models. Further, the separation of tariffing flexibility from the need for technology standardisation is one of the goals of the project.

2.3.5 Security versus performance

It is well known that security requirements can lead to reduced performance. Therefore, wherever possible, we should use 'structural security' rather than cryptography or per access checks. An example of structural security is random audit, where customers are trusted to behave because of a broader deterrent.

2.3.6 Performance versus system flexibility

Flexibility is often provided by introduction of layers of abstraction. However, such layers often reduce performance. Nonetheless, good design can ensure this is not the case, where an abstraction is configured into place for instance by dynamic loading of compiled modules. Flexibility is more likely if, for each aspect of the system, a model that is the superset of all other models is identified - whether it be pricing frequency, pricing granularity, payment frequency, who pays, how they pay, when they pay, who trusts whom, who needs accounting information, how often they need it, what is being measured and so on. However, generality does not imply abstraction - these general models must be translated into real mechanisms. Solutions that give flexibility without impacting performance are further goals of the project.

2.3.7 Manageability versus system flexibility

A high level goal for our system is as much self-management as possible. Often manual management is used to provide capabilities the system cannot provide as it stands. Therefore, these two requirements stand implacably against each other. One can only claim that the self-managed functions provided are sufficient to reduce the need for manual intervention. However, this is difficult to measure or refute.

2.3.8 Security versus ease of use

A general approach stated as a goal in this project's definition, is to allow customers to "just do it". That is, to use the network in arbitrary ways without pre-agreement or planning. The caveat is "as long as they are prepared to pay". Thus, the solution to this dilemma is to introduce a single credit check and ensure the customer is always identifiable, then she may do anything that costs less than her creditworthiness without further need to request any permissions, as long as she has signed a contract agreeing to pay for the consequences of her actions. In practice, creditworthiness can be controlled by a variety of more specific models such as pre-payment or deposit, which allow for anonymity, as well as a more traditional telco account relationship. This is not to say that it should not be possible to offer services via a secondary access check, but it is not *necessary* in a market-controlled system unless it is *desired*.

3 Roles of Internet Stakeholders

In this chapter we present a role model that can be used to describe properties of stakeholders and essential relations between them. Considering that services to end users will not be provided by a single stakeholder but through a cooperative effort of several stakeholders, there is a need to understand how end users' requirements will be reflected in the relations between them. Also, even if in the preceding chapter it was found convenient to use everyday-language notions of "user", "provider" and "system" in formulation of the general requirements, it will, in order to operationalise requirements and make them specific, be necessary to more precisely specify which stakeholders there are and at which interfaces requirements apply.

3.1 Role model

A role model is comprised of roles and their relations. A role represents a group of functions, enabling anyone taking on the role to provide a set of services to its environment. It is essential to separate the role model from a particular configuration of an enterprise or a market place. Therefore, we define that an enterprise (named *stakeholder*) can take on several roles. Other stakeholders in the market place can take on a subset or the same set of roles (i.e. competitors are present in the same market place). There is no single unique mapping between the role model and a stakeholder.

In the forthcoming discussion on business models, it will be natural to think of roles as objects that can be defined by templates and instantiated in the description of specific scenarios. So, even if it is not considered necessary to make the scenario models overly formal, we start in this spirit by defining the general notion of a role.

3.1.1 Generic role definition

A role is a set of functions that conveniently can be considered to form a segregated unit. A role instance will interact co-operatively with other role instances in the process of service provisioning. For the sake of scenario modelling, it will be assumed that the functions of a role can be grouped into:

- Service functions
  A role instance will generally consume services from its environment and combine the consumed services with its own service function in order to again offer other services to the environment. To perform this operation, the role instance will use resources at its disposal.
- Charging functions
  Enabling the monitoring of resource usage and service deployment. A charging function can interact with similar functions in the environment. Such interactions would typically be reports on resource usage in some agreed format. If enough of the internal resource usage of a third party is visible within the service function associated with a given charging function, performing charging for such third parties can be offered as a service in itself.
- Business policy function
  The business policy of a role instance are the decisions taken on (1) services, tariffs and possible service level agreements associated with services to be offered to the environment, and (2) how to deploy services offered to the role instance based on their prices, service characteristics and possible associated service level

agreements. The business policy function takes internal input from the service and charging functions. External exchanges are tariffs, service definitions and service agreements; generally the information needed on whom should pay whom what for what service.

The relations between the functions of a role and the interfaces between a role and the environment can be pictured as follows:

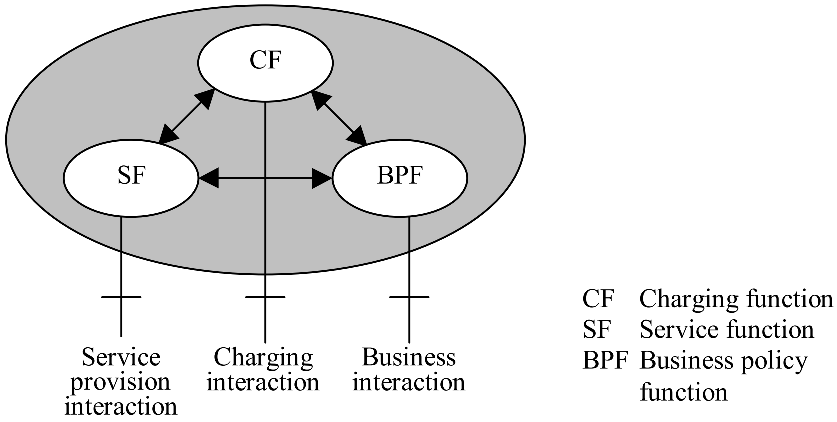

Figure 3.1 Role functions and interactions

Not all functions need to be present in all types of roles. Also a stakeholder taking on several roles will usually choose to organise charging in a centralised hierarchy, leaving only rudimentary charging functionality distributed. Likewise, a stakeholder can offer centralised charging functions to others through charging transactions.

While the service provision and charging interaction points are technical interfaces, the business interaction point may or may not have a technical counterpart.

3.1.2 Classification of roles

In order to distinguish stakeholders in more detail, we classify the stakeholders according to the type of service they provide or consume. We suggest introducing two service layers reflecting two main types of services: the ***Infrastructure Layer*** and the ***Internet Service Layer*** (see figure below).

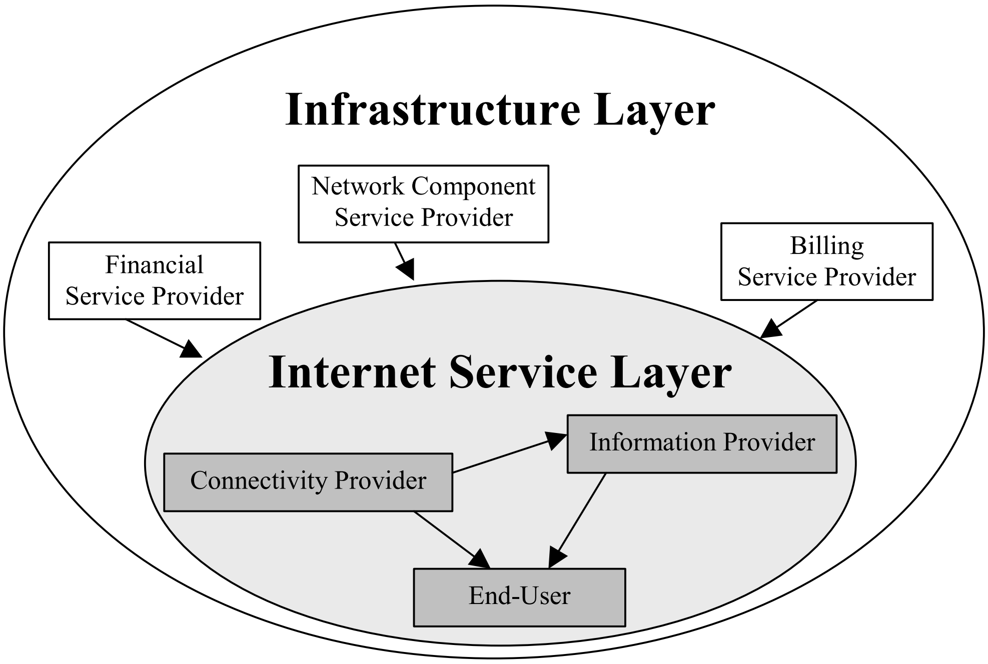

Figure 3.2 Service Provider Model

3.2 Internet Service Layer

The stakeholders of the Internet service layer are characterised by providing and consuming Internet services. An ***Internet Service*** is defined as a service that can be provided on the Internet. At this level of abstraction, the Internet service can be a network service (e.g. IP service, RSVP service) as well as an information service (e.g. stock market quotes). The stakeholders at the Internet Service Layer can be grouped into three classes: ***Connectivity Provider***, ***Information Provider***, and ***End-User***. Stakeholders who belong to the connectivity provider or information provider, are also called ***Internet Service Provider***:

- ***Connectivity Provider***: A connectivity provider is a stakeholder who provides the means to forward IP data packets on its network. A connectivity provider can be an access provider, a backbone provider, a server farm provider, and an end-user network provider.
- ***Information Provider***: The stakeholder in this role provides services on top of the network services provided by the connectivity provider. An information provider processes and supplies information to consumers. Application service providers, content providers, Internet retailer, communication service providers, or market place provider belong to this class of service provider.
- ***End-User***: In contrast to the definition of a consumer, an end-user is a consumer who does not resell a service, which it consumes. A device or application can also consume the service on behalf of the end-user.

In order to describe the business relationship between information providers, connectivity providers, and end-user, a further subdivision of these stakeholders is helpful.

3.2.1 Connectivity Provider

The connectivity providers manage networks and computers. These providers can be classified according to the functionality of the network within the Internet. Therefore, we can distinguish four types of connectivity provider:

- ***End-User Network Provider***: The stakeholder in this role can be the end-user itself or a corporation. The end-user is responsible for the network (e.g. a single PC or a LAN) or uses the network of a corporation she belongs to.
- ***Access Provider***: An access provider covers the 'last mile' between the end user and the backbone provider, utilizing copper lines, fiber line, or radio technologies. The dial-in modem provider (e.g. AOL) and the local telephone companies (e.g. Pacific Bell or SBC), that provide the telephone line for connecting to the Internet, are access providers. Other examples of access providers are mobile service providers (e.g. Vodafone) or wire-less service providers (e.g. @speed).
- ***Backbone Provider***: A backbone provider connects access providers to its high capacity network. Examples for backbone providers are AT&T, MCI Worldcom, British Telecom, Telenor/Nextra, Global Crossing, Qwest, and Level3.
- ***Server farm Provider***: The role of this stakeholder is to provide a secure computing facility to information provider, guaranteeing high reliability and availability of their servers and high-speed connectivity to backbone providers. Examples of this kind of service providers are Exodus and Akamai.

3.2.2 Information Provider

The information provider can be classified according to the kind of information provided. We distinguish five different roles that a stakeholder can take on:

- ***Application Service Provider***: The service of an application service provider comprises the lease of usage time of software applications they own. The application software provider takes care of maintenance and management of the software. Examples of such applications are www-server (web hosting), SAP and FileMaker.
- ***Content Provider***: Content providers collect, organize, and present information, which is available on the Internet. There are content providers specializing on certain topics (e.g. Marketwatch, CNN) or help people to access information quicker (e.g. Yahoo).
- ***Internet Retailer***: A stakeholder in the role of an Internet retailer sells products on the Internet. Examples for product retailers are Amazon.com and Barnesandnoble.com.
- ***Communication Service Provider***: The communication service provider offers services like Internet Telephony, email, or fax. Companies like Net2Phone, AOL, and efax.com belong to this group of Internet service providers.
- ***Market Place Provider***: A market place provider either brokers information of other service providers or provides an environment for providers to offer their services. In both cases, consumers can easily evaluate services. Examples of this kind of service provider are brokers and electronic market places like BandX.

3.3 Infrastructure Layer

The infrastructure layer consists of all those providers that provide service to stakeholders of the Internet service layer, but are not directly involved in the service provision between the stakeholders in the Internet service layer. The service provided may or may not use the Internet. Some of the stakeholders that belong to the Infrastructure layer are:

- ***Network Component Provider***: A stakeholder in the network component provider role owns network lines or computers, which it leases to other stakeholders.
- ***Financial Service Provider***: A financial service provider provides a service for completing the money transfer. An example for such a stakeholder is a credit card company.
- ***Billing Service Provider***: This stakeholder's role is to provide billing services to stakeholders who outsource their billing of customers

3.4 Business Relationship Model

The business relationship model describes the relationships between the Internet stakeholders. The basic business relationship between the connectivity provider, the information provider, and the end-user is presented in Figure 3.3. The black arrows represent the direction of the service delivery: $ISP_x \rightarrow ISP_y$ means *ISP_x delivers service to ISP_y*. The connectivity provider might provide services to the information provider, the end-user, and to the connectivity provider. The information providers only offer services to end-user and other information provider. The end-user only consumes services.

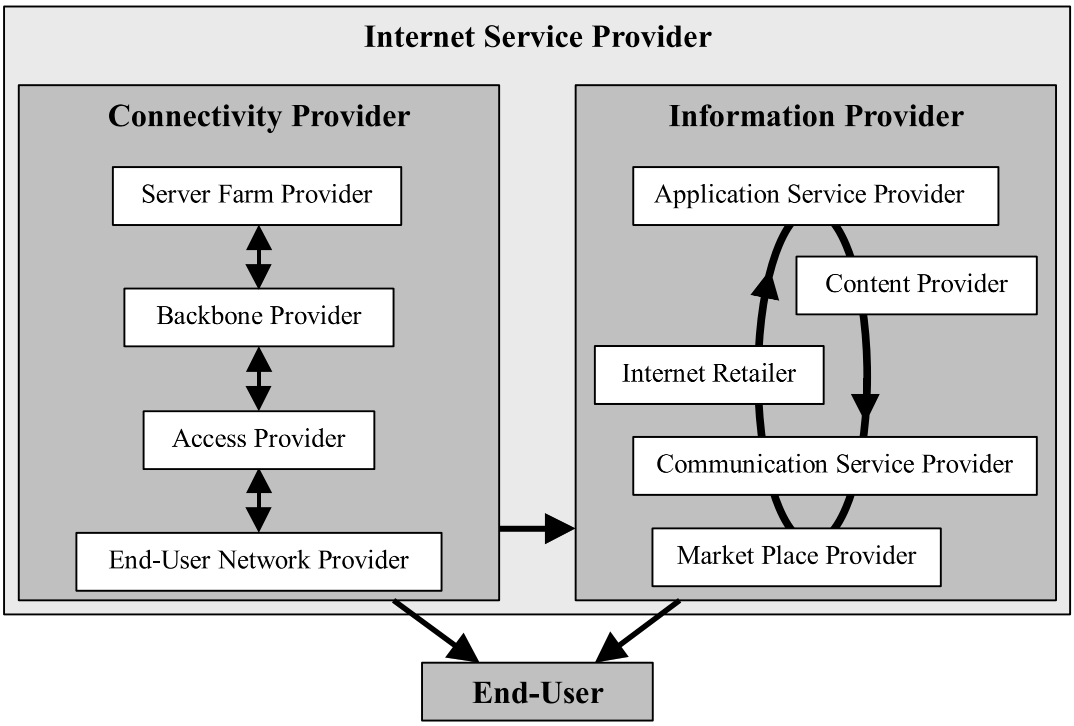

Figure 3.3 Internet Market Layer

The interaction among the connectivity providers is determined by their functionality in the Internet. The end-user's network is always connected to the access provider's network, whereas the access provider has at least one connection to the backbone. The server farm providers are located within the backbone. (see black arrows within the connectivity provider box in Figure 3.3).

The information providers are more flexible with regard to their business relationship among each other. All kinds of relationships are possible (represented by the circle in Figure 3.3).

As an exercise in using the framework provided in the previous sections, it will be applied to an example scenario of current Internet service provisioning. Focus will be on describing the service provisioning- and business relationships of a relatively simple Internet service provider. In this example scenario, the ISP will encompass the roles of a backbone provider and a communication service provider. In the former respect its specific tasks will be to provide the physical access point to the Internet (e.g. a modem), the internet address to the customer and the necessary authentication and authorisation associated with giving users access, and to run a network of several IP-nodes connecting access points at different locations. In order to operate the IP backbone the ISP owns a set of routers connected through leased lines. As a communication service provider the ISP will offer its customers access to a mail service and a web hosting service.

Another stakeholder required in the scenario is an infrastructure operator, taking on the Access Provider and Backbone Provider roles. The infrastructure operator is also a telephone network operator, and having a customer relationship with an ISP presupposes in this example a customer relationship to an infrastructure/telephony operator. Whereas the ISP runs an IP-backbone, the infrastructure operator runs an infrastructure backbone offering leased line services. The Infrastructure operator also operates a charging function, which makes possible usage based charging of the access to the ISP.

The scenario is depicted in Figure 3.4. As may be observed, there are altogether five stakeholders present. In addition to the customer, the ISP and the Infrastructure/telephony operator, the scenario includes a separate application/content provider stakeholder, and a second Internet backbone provider that may be considered as an upstream connection to the global Internet. Keeping in mind that the main focus is on the ISP stakeholder, interface relations between this and the stakeholders towards which it has direct relationships are included, denoted I1 – I5 in the figure.

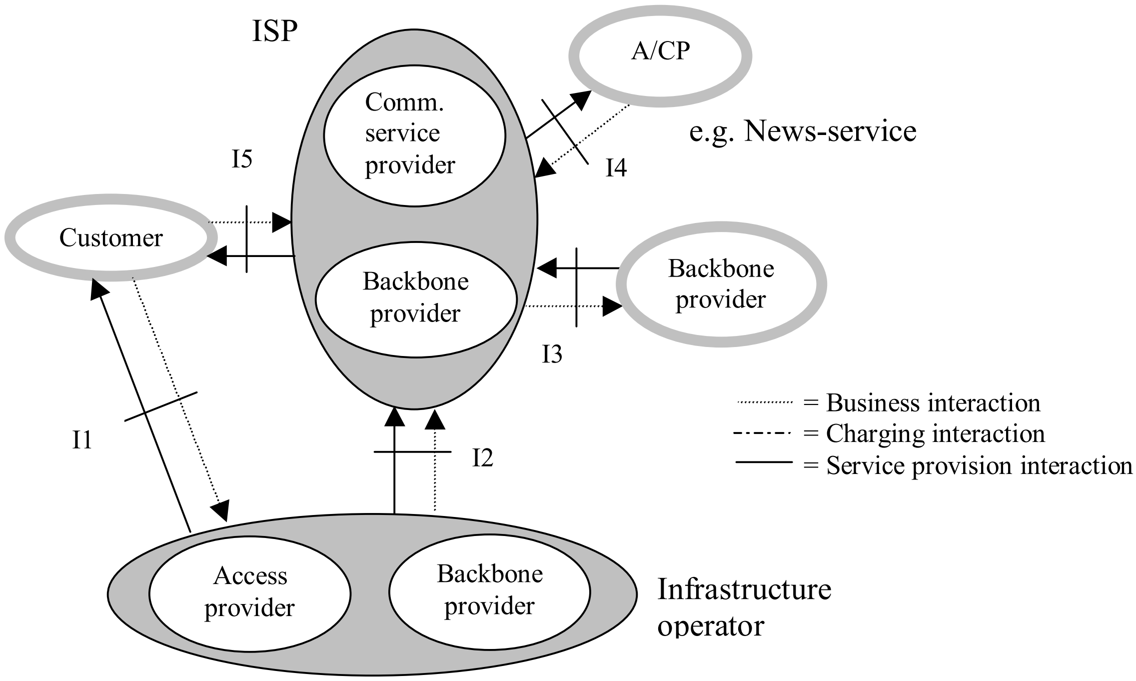

Figure 3.4 Example ISP scenario

In accordance with the general role definition introduced, each interface has three aspects: service interactions, charging interactions and business interactions. In this particular example, there are no direct interactions between charging functions of stakeholders, hence no charging interaction lines are depicted in the figure. The semantic of the arrows is that for service provisioning interactions, arrows point in the direction of service provisioning (which might be opposite to main flow of information), and for business interactions arrows point in the direction of net value flow. The five interfaces may be (summarily) described as follows:

Interface I1 (Customer - Infrastructure operator)

- **Service provisioning**
  The Infrastructure operator offers a switched analogue (telephony) or digital (ISDN) communication service.
- **Business interaction**
  Tariffs offered are three part tariffs consisting of a fixed monthly fee, per minute connection charges and per connection setup charges. Customers get information on usage and corresponding charges through quarterly billing.

Interface I2 (ISP – Infrastructure operator)

- **Service provisioning**
  The infrastructure operator connects incoming calls to ISPs access points (AP). These are known to customers through a single address (telephone number). The infrastructure operator performs intelligent routing in order to utilise its own network resources and the ISP's line termination resources. Also, the infrastructure operator provides the leased lines interconnecting the ISP's nodes.

- **Business interaction**
  The ISP has a termination agreement with the infrastructure operator, and will thus get part of the usage based income associated with calls to its APs. For the leased lines, tariffs are based on capacity and distance.

Interface I3 (ISP – other backbone provider)

- **Service provisioning**
  The service is a best-effort transfer of datagrams over a higher level in the Internet routing hierarchy. The other backbone provider thus provides global connectivity. In addition to datagram transfers, interactions take place in the global distributed domain-name directory service.
- **Business interaction**
  The ISP pays (in this example) a flat fee for the global Internet access. Service level agreements on availability and capacity might apply.

Interface I4 (ISP – other application/content provider)

- **Service provisioning**
  The ISP provides a best-effort datagram service.
- **Business interaction**
  The A/CP pays a flat monthly fee dependent on the access capacity.

Interface I5 (ISP – Customer)

- **Service provisioning**
  At the network level the ISP provides a best-effort datagram service. As access is dial-up, customers are authenticated and given service authorisations at connect time. Also, the customers are supplied with IP-addresses as they enter the service. Customers are given access to the domain name directory service. At the application level, customers are provided with a mail service and a web hosting service.
- **Business Interaction**
  Customers pay a flat monthly fee.

The main concerns in this context are service provisioning, service quality and charging, so the system requirements relevant to sustain the above business scenario are:

For the infrastructure operator

- Collect duration charging information for connections between I1 and I2, monitor connection establishment and information transfer in order to relate to overall QoS objectives (e.g. in terms of call blocking)

For ISP

- Monitor connection usage at I2 in order to collate accounting with Infrastructure operator
- Monitor modem pool blocking (possibly with the aid of the infrastructure operator) in order to relate to overall QoS objectives
- Monitor IP traffic at I3 and I4 in order to decide on possible terms of connection agreements.

Vital parts of the business policy function of the ISP are tariffs for various access alternatives, dimensioning of lines leased from the Infrastructure operator going into its IP-backbone, and the interconnect capacity to the global Internet.

The above example is aimed at being relevant in the current marketplace. Considering however the diversity of the industry, it is difficult to create a scenario that can be claimed to be representative. A synopsis of various current stakeholders is thus presented in the next section.

3.5 Examples of current ISP business models

In order to better understand the environment in which the ideas of the M3I-project will have to survive, descriptions of six stakeholders currently observed in the marketplace are given.

It is difficult to determine whether and how much money is transferred between two actual stakeholders in a business relationship and what other conditions apply in terms of charging and service provision. The direction of the money flow and the amount depends mainly on the market power of the ISP, i.e. the size of the network, the number of end-users subscribed, and the kind of service provided. Therefore, the examples of business relationships given in the following subsections do not contain a detailed analysis of the various types of interactions previously identified. However, in order to get a better understanding, we give some instances of possible revenue flows here:

- Backbone providers may get usage-based revenue from termination agreements with access providers: Access providers charge customers according to the time being connected to the network, and this charge is split between the access provider and the backbone provider.
- Revenue of an ISP with many subscribers might also come from agreements with application providers or content providers connected directly to the ISPs network.
- Backbone providers' revenue streams depend on the size of the network, the location of the network within the global network, and the nature of the information flow over the interconnection point.

Beside the description of the business model of each company within the subsection, we also illustrate their business relationships in the terminology of the reference model. All data presented was gathered on May 1st, 2000.

3.5.1 Mindspring / Earthlink

Mindspring is located in the USA and offers service in the USA and the United Kingdom. Its market capitalization is $2400 million. The annualized revenues of $1300 million are expected from 4.2 million subscriber accounts and 148000 web hosting accounts.

Mindspring/Earthlink offers web-site hosting and Internet access service. The network used is leased from Sprint and PSInet. Only the dial-up access sites in southern California are owned and operated by Mindspring / Earthlink.

The Internet access services cover a wide range of access technologies (e.g. dial-up modem, Frame-Relay, ISDN, or DSL) and access speeds. The price for modem dial-up service depends on the number of hours that a user wants to use the service. The prices on May 1st, 2000 ranged from $6.95 for 5 hour/month (each additional hour costs

$2) to $26.95 for unlimited access. This service is bundled with web space and a certain number of mailboxes.

The web-site hosting service provides an environment for enterprises to run a web-site. The service comprises a fast reliable web server, POP email, CGI scripts, web statistics reporting, storage, an allowance of a certain amount of traffic to the web-site (2Gbyte-6Gbyte; all additional traffic is charged at a rate of 0.04 – 0.1 $/Mbyte), and an e-commerce software support. The web hosting service is actually split-up into four sub-services in order to address the need of different enterprises.

Regarding Mindspring / Earthlink’s business relationship with the end-user, the company takes on the roles as an access service provider and a communication service provider. The company is in the role of a communication service provider if it leases the backbone or access network.

3.5.2 Exodus

Exodus focuses on the Internet server farm market and owns a world-wide backbone. Its server farms are located in Asia, Australia, North America, and Europe (Sweden, United Kingdom, Germany, Netherlands, France). The annual revenues are 242 million and the market capitalization is $16900 million.

Its service comprise the hosting of servers in a highly secure location, featuring raised floors, temperature control system, seismically braced racks, smoke detection, fire suppression system, motion sensors, and surveillance cameras. In addition to this, there are redundant power supplies, multiple backup power generators, and multiple fiber trunks coming into the server farm. The connection speed between the customer’s servers at the server farm and the server farm network can be 10Mbps Ethernet, 100Mbps Fast Ethernet, or 1Gbps Ethernet. Customers can also get multiple LAN connection in order to be fault tolerant against network connection failures. Exodus charges either a flat rate or a usage-based charge for those connections. In order to analyze the utilization of the connections, Exodus provides customers with a bandwidth report, containing detailed information about the line usage. In addition, it provides connectivity reports to the main ISPs and route information of IP packets.

Beside the server farm, Exodus owns a backbone that connects all server farms redundantly, allowing single point-to-point fiber link failures. The company also offers Internet access services at T1 speed (1.54Mbps) and DS3 speed (45Mbps) for a flat rate or usage-based rate. Especially, they connect the customer’s offices with the server farms.

With regard to Exodus’ business with enterprises, the company is in the role of an access provider. However, their main business is their server farm and backbone service.

3.5.3 America On-Line

AOL is a world-wide operating corporation, providing service in the United Kingdom (service name: AOL, Netscape Online), Germany (service name: AOL and CompuServe), and France (service name: AOL). The total customer base is 22 million. The current market capitalization is $132000 million and the revenues are approximately $4700 million.

AOL also bundles Internet connectivity services and information services. The connectivity service is a resale of connectivity service of Sprint, GTE Internetworking (formerly BBN Corporation), and MCI Worldcom Advanced Networks. AOL itself does not own any networks. Instead, AOL concentrates its efforts on providing content, retailing, and communication services to their customers. AOL has more than 1000 e-commerce and content partners. In many cases, AOL owns a stake in the partner's company. The company also wholly owns Digital City Inc (i.e. local content service provider), ICQ (i.e. chat communication service provider), MovieFone Inc (i.e. movie content service provider), and Spinner.com (i.e. music content service provider). That means, revenues of AOL come from the resale of the connectivity service, the advertisements placed on the AOL and Netscape Netcenter portal sites, and their communication service tools (e.g. AOL Instant Messenger, Net2Phone).

The roles played by AOL are access provider, backbone provider, communication service provider, Internet retailer, and content provider, considering AOL's business relationship with the end-user. The business relationship with connectivity providers puts AOL in the role of an information provider.

3.5.4 RedNet

RedNet is an Internet service provider that is privately owned and solely available in the United Kingdom.

The dial-up access service comprises 10 e-mail addresses, 10 Mbytes disk space for WWW site, and customer support. The service plan for the remote access service is a flat rate for Internet service in addition to the local call rate charged by the local telcos in the UK. An extra charge applies for domain name service, email hosting, increased disk space for the web site, ftp disk space, own domain name web site, and secure transaction over the www.

In addition to this, RedNet is offering services focusing on business customer. Those services comprise firewalls, VPN, leased Internet access lines, ISDN, Web-site design, email gateways, application hosting (e.g. FileMaker pro), and server hosting. The company offers four types of server hosting (bronze, Silver, gold, and platinum). The bronze service starts with providing a dedicated line to the Internet. The silver service comprises the fault tolerance mechanism like backup of data and data recovery after server crash. The gold service gives the enterprise support in optimizing the server performance. The difference between the gold service and the platinum service is that the customers of the platinum service are allowed to use streaming audio and video on their web-site.

RedNet's roles in the business relationship with end-users are an access provider, backbone provider, server farm provider, application service provider, and communication service provider. RedNet is only in the role of an access provider and backbone provider, if the company sells it services to other connectivity provider.

3.5.5 Covad

Covad currently offers its Internet access services only in the USA, focusing on the major metropolitan areas. The revenues are $66 million and the market capitalization is $4189 million.

This company owns their own network, which includes an USA-wide backbone and the hardware equipment collocated at the ILEC central offices.

As of May 1st, 2000, Covad offers two DSL services, targeted at residential users and businesses. While the basic service comprises two choices (608/128kbps and

1500/384kbps), the business-targeted service offers six different choices between 144/144kbps and 1.1/1.1Mbps. The business services are further differentiated with regard to the number of users supported. In addition to this, Covad offers VPN service between company branches and headquarters as well as between company network and employees. The data between the Covad network and the company network is encrypted. However, Covad does not sell its products directly to end-users. The companies customers are Internet service provider like Mindspring / Earthlink.

In addition to this, Covad owns the subsidiary LaserLink.net that offers services to set up a virtual Internet service provider. Covad manages and administrates the network while the customer sells the service as its own (wholesale business).

The roles of the company in the business relationship to information provider are an access provider and backbone provider. In addition to this, Covad takes on roles as an access provider, a backbone provider, and a communication service provider in its wholesale business.

3.5.6 Akamai

Akamai is an Internet service provider that offers services to content provider guaranteeing fast and reliable delivery of content to end-users. The company is based in the USA but operates in

In order to be able to offer such a reliable and fast service, Akamai distributed 4000 cache servers at the edge of the Internet world-wide addressing two issues. First, routing of content via various paths between their servers becomes possible even if certain parts of the network are congested or out of service. To detect such network failures, real-time information about the network is constantly analyzed by fault tolerance software. Second, servers are always physically closed to end-users, resulting in short responds times to end-users content requests and high bandwidth availability to the users computer. Content has not to be transmitted across the several backbones without any control of the quality of service.

Especially, Akamai focuses on certain content type. It offers streaming media service, including live events, continuous broadcasts, and on-demand media. Beside their server technology, Akamai uses forward error correcting software and multiple copies of streams to guarantee the delivery of the content to all their servers at the edges of the Internet.

Akamai has business relationships with content providers, backbone provider, and access providers. The company has to rent space for their cache servers in the central offices of the access provider and connect them to the network of the access provider. It also has to connect their servers to several backbone networks. The service agreement with content providers has to specify the kind of content, which has to be cached. Akamai's roles are a server farm provider and an application service provider. The company takes on the role of a server farm provider since it owns cache servers. Since the company runs software to distribute the content, it is also in the role of an application service provider.

4 Specific role relationships

In the ensuing sections we look at some special relations between stakeholders involving interactions between the business policy functions of their roles. These interactions fulfil some important needs in an Internet with various yet unspecified pricing models, and various interconnect relations between stakeholders. By integrating these interactions into the previously described roles, a number of new stakeholders in the Internet market may arise, whose business is to simplify underlying internet service prices and resell them in some sense. We describe a number of cases.

Section 4.1 describes a “clearinghouse” whose function is to gather charges charged to all end-customers of some communication, and to redistribute those charges according to some agreement among the end-customers. This function is also needed in creating end-to-end prices and bundling connectivity charges with information charges.

Section 4.2 describes a “dynamic price handler” whose function is mainly to automate a single end-customer’s reactions to dynamic connectivity charges, according to some customer strategy.

Section 4.3 describes a “risk broker” whose function is to create a “few” transport services with “simple” end-to-end prices, from underlying connectivity provider services with highly varying prices or service qualities.

The new functions described are still generic in the sense that the underlying services and tariff models are not specified, nor the nature of interconnect agreements. Nevertheless a number of quite generic requirements can be made on the environment of these functions. These requirements carry over to any stakeholder integrating these functions, so it is natural to state them in a general context.

4.1 Clearinghouse

IP-sessions may generate two or more bills, one for the sending end, one for the receiving end. The task of creating an IP-session with *one* price, which may be bundled with an application or which may be distributed in arbitrary manner among the end-customers, is the function of “clearing”. This function may be performed by a new stakeholder in the market (a *clearinghouse*), or integrated into one of the existing roles. The clearinghouse task can be viewed as a special aspect of the business policy function generally associated with a role.

The action of clearing involves two or more paying parties and two or more parties (typically the end-customers’ Connectivity Providers) demanding payment of a certain amount by the customers. The clearinghouse collects the demands from the parties to be paid, sums the demands, collects the percentages of the total amount each customer is willing to pay, and announces to each paying party the amount to pay. This process is depicted in Figure 4.1.

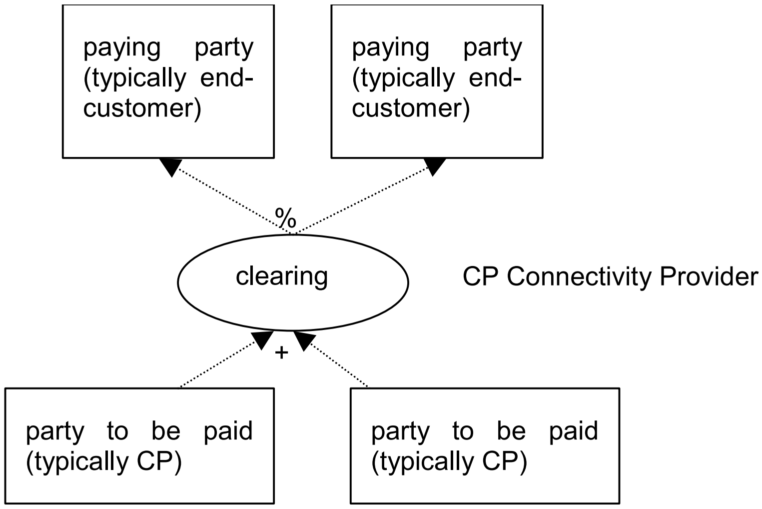

Figure 4.1: Clearinghouse with flow of demands for payment

There may be more than two paying parties, and there may be more than two parties to be paid.

The clearinghouse sums up incoming charges and distributes the sum to the paying parties, *after* communication has taken place or on-line, during communication. The clearinghouse expects all paying parties to pay the agreed share of the total cost of communication without repudiation. Note that the total cost of communication may not be known beforehand.

A clearinghouse in its basic form does not take on the responsibility for the services provided to the paying parties. I might choose to do so, but will then take on the role as a service provider towards the paying parties. This will have implications for the business policy function of that stakeholder in excess of what is required for clearing.

A natural function of the clearinghouse is the distribution of payment back to the parties to be paid. This means the parties involved need to agree on a common mode of payment: exchangeability of currencies, agreement on set of credit cards or form of micro-payments. A clearinghouse may also typically reduce the cost of multiple transactions, by batching up funds transfers to its regular customers and suppliers (the Connectivity Providers and big Information Providers).

4.1.1 Requirements on clearinghouse's environment

- The various components of the service for which the charges should be combined must be possible to synchronise. The need for synchroniseable service elements is not generally catered for in the connectionless Internet paradigm. It might sometimes ultimately be impossible to intercept connectivity usage charges associated with a given combination of application/connectivity services. A solution to this could be to employ charging functions at the end-points (e.g. with the customer and application service provider) in order to perform end-to-end resolvement of total charge.
- If the clearing involves redirection of charges relative to some default arrangement, necessary information to perform the redirection must be available to the charging party at service activation.
- The agreement procedure for charge distribution between paying parties needs to be specified. Such procedures might be application specific, but the general requirements associated with such procedures still need to be understood.

- The clearinghouse should be able to follow up arguments about incurred charges (but not service itself). In order to achieve this, the clearinghouse may need access to accounting records belonging to the identified service components.
- The clearinghouse takes no responsibility for the quality or the provision of the service. The clearinghouse does not modify the information flow. The clearinghouse expects all paying parties to have agreed on a service and service level beforehand. It expects the involved providers to deliver the service at appropriate level, and it does not participate in arguments about service level.
- The involved parties need to have agreed on a common mode of payment (currency, credit card). It is out of scope for M3I to define e-commerce procedures, but we should keep the requirement in mind, as it may restrict the scope of clearing.

It may seem more practical to have the clearinghouse declare the total bill *before* the service has been delivered. In this case the clearinghouse takes the risk of paying the connectivity providers (parties to be paid) independently of what it may receive from the paying parties. This is a new functionality, risk brokering, which is described in more detail in the later section on the “risk broker”.

4.1.2 Application of clearinghouse in some interconnect environments

Applications of a clearinghouse are in

- “receiver pays call” independently of whether intermediate connectivity providers charge each other after the “receiver pays” principle
- achieving any desired kind of charge distribution between participants in a multicast tree independently of charging principles supported by intermediate connectivity providers
- Bundling all arising connectivity provision charges with the application charge

Clearinghouse summing and distributing local costs

Assume that the connectivity providers involved in a communication expect to get paid “local costs” by one or several of the communication ends. Clearing may help to achieve this, as illustrated in Figure 4.2.

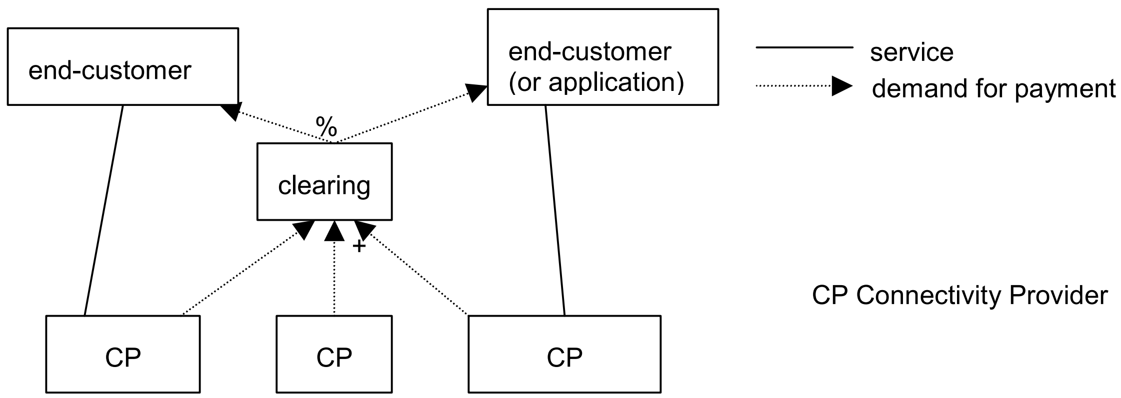

Figure 4.2: Clearing (distributed or not) among all connectivity providers in the communication

The clearing function may architecturally be realised by several components distributed among the connectivity providers which efficiently work together to collect the total demand for payments, distribute the demands in the right percentage to the paying parties, and in some non-specified way distribute the actual payment to the connectivity providers involved in communication. An efficient signalling mechanism achieving the clearinghouse functionality is described in [4]. The depicted clearing function may generate a bill that specifies all incurred local costs. In whichever way the clearinghouse is being implemented, one has to keep in mind the requirements made about the clearinghouse (Section 4.1.1):

- The requirement on common mode of payment may severely restrict the extent of “clearing integrated with connectivity provider-service”.
- The requirement on identifiability of communication to be paid for seems more easy to meet in connection-oriented solutions (RSVP, e.g.).
- Responsibility for quality of service together with price level is achieved by integrating clearing into connection set-up, as in [4].

Clearinghouse in an edge-pricing situation

Suppose connectivity providers charge each other only for flows crossing a boundary in between, and that connectivity providers charge end-customers for flows crossing some boundary (however flow is defined). The total charge of a particular communication between sender and receiver must then only be collected from the sender’s and the receiver’s connectivity provider.

The demands for payment in a “receiver-pays call” are illustrated below.

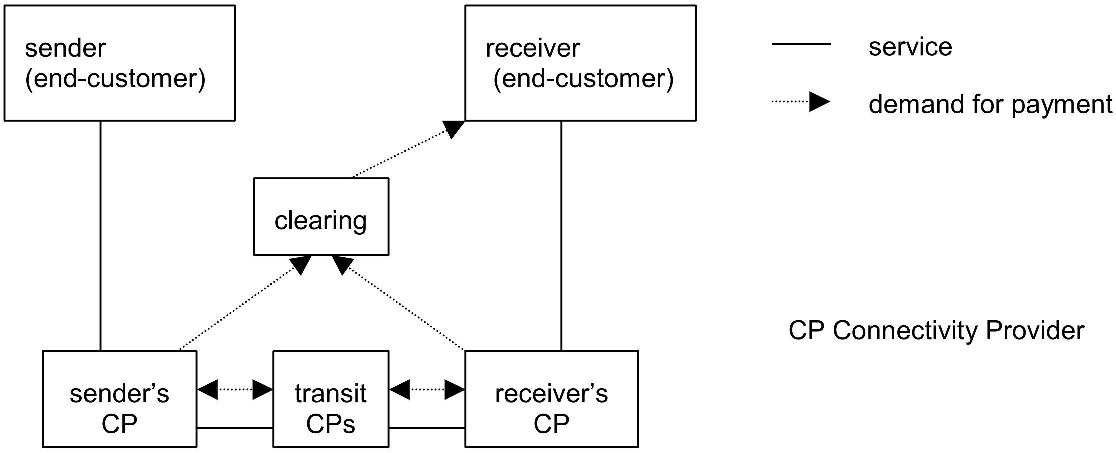

Figure 4.3: Flow of demands for payment in a “receiver pays call” for edge-pricing situation

The double arrows in the figure illustrate that demands for payment are accounted for between connectivity providers.

In this scenario the clearing function may belong to an stakeholder other than the involved connectivity providers. Moreover, clearing is not necessarily coupled to a concrete implementation of the service (connection set-up, e.g.). A clearinghouse for this scenario is described in [1].

The requirement of “common mode of payment” is easier resolved in this scenario than in the previous, since this scenario involves fewer parties.

The requirement on identifiability of flow belonging to the particular communication has to be investigated further, as well as the issue of responsibility for quality-of-service.

4.2 Dynamic Price Handler

Certain pricing models operate with frequent price changes on short time scales in order to achieve "optimal" load management. An implicit economic assumption behind these models is that customers react to the price changes at the same time scale. In order to relieve the (human) customers from the burden of frequently reacting to price changes, parts of the customers' service policy function can be automated. We describe this special aspect of a customer service policy function, the main assumptions, and possible applications in business models.

The dynamic price handler has a relationship to three roles, typically a connectivity provider, a (human) end-customer and an application provider. It transforms dynamically varying connectivity provider-charges into a total charge by dynamically varying service classes according to the end-customer's strategy.

The term "variation of service class" may mean many different things depending on the price and service model. E.g. choosing a diffserv type-of-service, a priority level, the bid in an auction, or simply transmitting at another bandwidth.

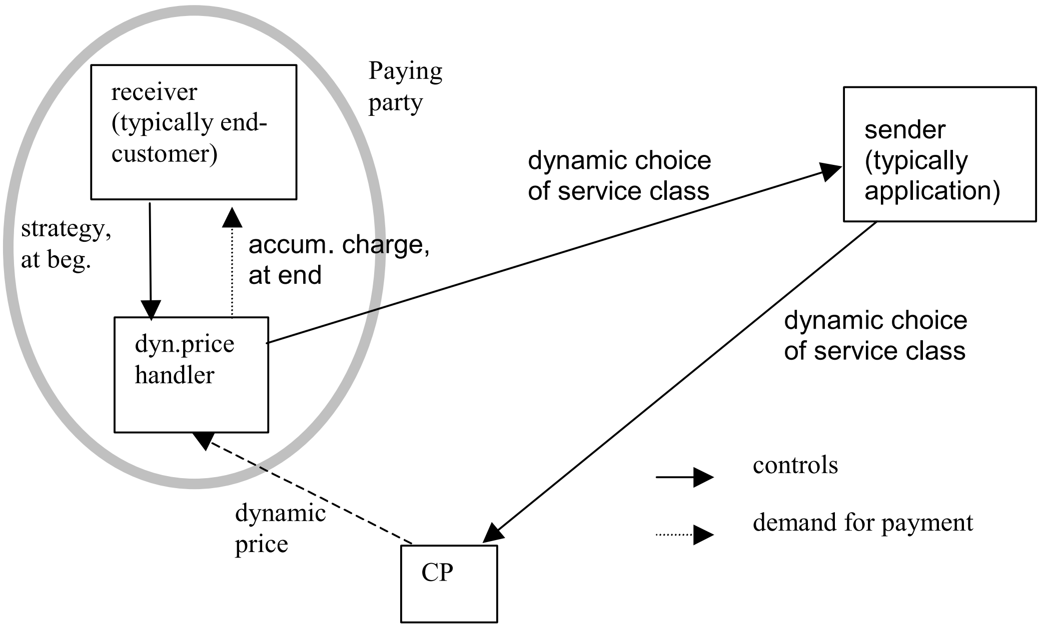

Figure 4.4: Dynamic price handler, simple example

To begin with a simple example, assume

- that all communication ends are located in the same connectivity providers domain,
- that the connectivity provider varies prices on a time scales shorter than the communication duration,
- that the sender ultimately controls the choice of service class,
- that the connectivity provider charges the receiver, and
- that it is the receiver who actually wishes to pay.

Then the receiver may wish to vary the service class during the session in response to the price signals. This may be done with the help of the dynamic price handler, depicted in Figure 4.4. In the example, the receiver's dynamic price handler sends control signals to the sender, who acts accordingly. A trust relationship must exist between the sender and receiver, as the receiver must trust the sender to act according to control signals sent.

More generally, the dynamic price handler acts on behalf of a paying party (the receiver in the simple example). This means that it expects a strategy from this party at the beginning of the session and it outputs a total charge towards the paying party at the end of the session. Examples of strategies are:

"minimise transfer time inside total budget,"

"disconnect when per-minute-charge higher than ...".

On a shorter time scale the dynamic price handler varies service classes in response to price.

Several modes of interaction between parties' service policy functions in relation to the existence of a dynamic price handler can be foreseen, depending on the environmental conditions of trust and competition. Some of these are:

- The paying party may have complete control both over the service and the dynamic price handler (no trust needed)
- The paying party can control the dynamic price handler and instruct the other party who controls the service (the example above)
- The paying party may control nothing. This is possible in a larger environment of competition or trust.
- The paying party may input a “high level” policy to a dynamic price handler controlled by the other party.

In a specific scenario involving a dynamic price handler, such interactions need to be defined.

4.2.1 Requirements on the dynamic price handler's environment

General requirements in connection with use of dynamic price handlers are:

- If more than one dynamic price handler is involved in a communication, possible conflicts in strategy and service control must be resolved.
- The exchange of control information between parties must be fast enough to facilitate useful reaction on price and service level signals.

4.3 Risk broker

In an environment with a network infrastructure that offers no QoS/price guarantees, but where users nevertheless have the possibility to dynamically adapt to price or quality signals, it may still be possible to offer service and or price guarantees to users by introducing brokering at the network edges. The risk broker is part of the service policy function of an stakeholder performing this task. The role of such an stakeholder involves risk brokering and clearing, and it interfaces one or more paying parties (typically end-customers), one or more parties to be paid (typically connectivity providers generating dynamic prices), and possibly an application/content provider.

The risk broker function has certain similarities to the dynamic price handler. While the dynamic price handler exist in an environment where end users or content providers relate directly to dynamic pricing, the risk broker function does however take part in the service provision of mediating between more traditional services and dynamically priced network services. The differences to the dynamic price handler are:

- The risk broker offers a list of transport services and attached charge (or a tariff, depending on the service) to its customers before or at the *beginning* of the session, at its own risk. This charge is to be paid by the paying party.
- The strategy that is input by the paying party consists solely in the choice of a transport service. The services provided must be defined somehow. One possibility is to use a well-established regime like Int-serv.
- The risk broker creates the agreed transport service by making appropriate use of the underlying dynamically priced connectivity provider service.
- The risk broker pays the connectivity provider its accumulated charges.

Note that the transport services offered by the QoS-provider need not necessarily be the same as the connectivity provider service classes. The QoS-provider may create transport services (with e.g. additional QoS guarantees) from underlying connectivity provider services. Risk brokering involves translation of usage incentives. The risks involved for the stakeholder performing risk brokering are fundamentally related to the differences in usage incentives offered and usage incentives faced.

The term “choice of connectivity provider service” may mean many different things, e.g. a diffserv type-of-service, a priority level, the bid in an auction, or simply transmitting at another bandwidth.

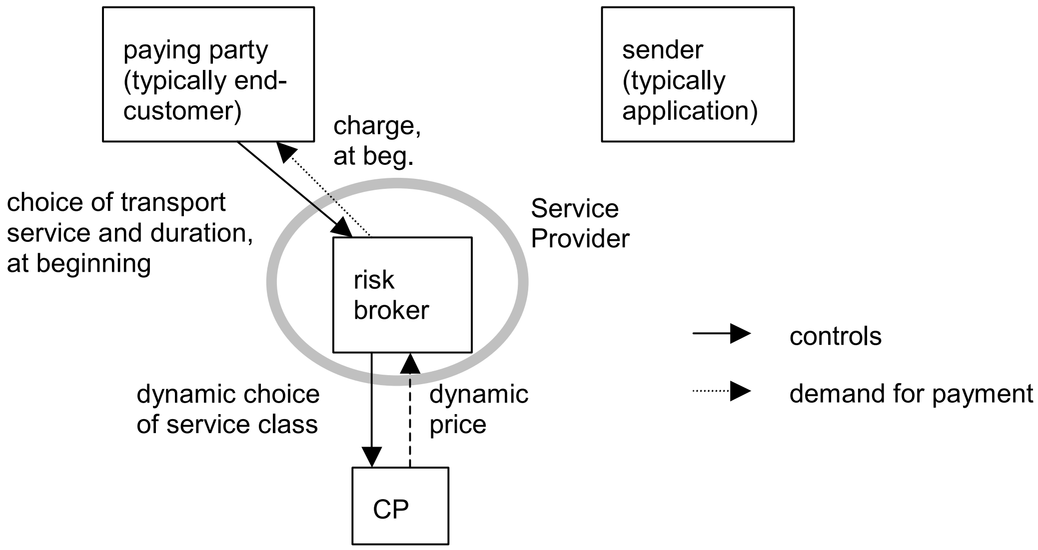

Figure 4.5: Risk broker

In analogy with the dynamic price handler, a simple example involving a risk broker will be given. The example is pictured in Figure 4.5. Only control and money flow interactions are shown in the figure. Assume

- a service provider that buys dynamically priced communication services from a connectivity provider and sells services with guarantees to third parties (“paying party” and “sender” in the figure),

- that this service provider is present at both points where traffic enters and leaves the underlying connectivity provider's network, and dynamically controls the choice of service classes towards it,
- that the receiver selects the service and takes the charge for it.

In this example, the service provider is a separate economic entity which competing edge will be the efficiency of its service policy function, i.e. the risk broker. The algorithms and possibly also the end-to-end protocols involved in the risk brokering will thus be business-critical for the service provider, and will probably not be generally available to third-parties.

4.3.1 Risk brokering and interworking

Several interconnect and interworking situations may be considered in connection with the existence of risk broker functions:

1. One service provider providing services over one underlying "dynamic pricing" connectivity provider. This is the situation most readily implemented, but also with the narrowest scope.
2. One service provider providing services over concatenated underlying "dynamic pricing" connectivity providers. In an edge-pricing paradigm, the underlying intermediate connectivity providers need not be aware of the existence of the service provider.
3. Interworking between service providers each of who again may provide services in modes 1 or 2 above or by more classical IntServ or DiffServ regimes. Services over the interconnect interface are the same as offered to end customers. This mode of interconnect requires the full set of QoS and interconnect apparatus under definition for interworking between telecom operators, see e.g. [23].
4. One might think of situations where the service provider offer services to a customer at one communication endpoint, but have a direct peer-to-peer interconnect arrangement with the customer at the other communication end. The distribution of responsibilities and liabilities for this setting requires further study.
5. A last possibility is where two service providers cooperate to provide services to end customers on an and-to-end and peer-to-peer basis. This setting is different from alternative 3 above in that the two providers must use compatible internal protocols and service policies. The distribution of responsibilities and liabilities will also here require further study.

4.3.2 Requirements on the risk broker's environment

General requirements in connection with use of a risk broker are:

- In interworking settings 4 and 5 above, a procedure for resolution of possible conflicting strategies must exist.
- The risk broker must trust or control the parties at communication end points not to make any choices of connectivity provider service classes on its own conflicting with the risk broker's. This implies that the risk broker needs to either monitor or intercept the packet stream.

5 Scenarios

The framework reference model end generic components described in the previous chapter are created in order to describe essential relations between stakeholders in a large variety of business configurations. Of this large number of possible configurations, two specific configurations have been selected in order to describe how, in a business model framework, requirements postulated in chapter 2 may be met. The descriptions will be made using the role model developed in chapter 3, including also the generic functions described and discussed in chapter 4. Relations between stakeholders will thus include specifications of how they interact with respect to service provision, charging and commercial aspects. A discussion of how the scenarios relate to the posed requirements follows in a later chapter.

5.1 The guaranteed stream provider

The motivation of this scenario is to provide a type of service to end-users that incorporates and extends the classical telephony like service. A guaranteed stream is applicable in circumstances where an increase in capacity to single users gives them a more than proportional increase in utility. In the overall optimisation of utility over the resource, some users will thus be blocked out while the remaining will be given capacities over critical thresholds in relation to their utility functions. Typical applications where such utility functions can be found are real-time audio and video services. In the current scenario, the users will not only be able to choose the capacities they want, they will also have the possibility to choose different qualities of service for a given capacity.

In M3I, it will be assumed that two stakeholders will cooperate in order to provide the above services to end-users: A stakeholder providing a basic communication mechanism and the other stakeholder making the refinement into guaranteed services. We consider a range of relations between the provider of the “basic service" and the “refined service”. At one end of the range, the basic service is a dynamically priced best effort service with no guarantees neither on quality nor price. At the other extreme, there is no refinement necessary, the basic provider delivers all relevant guarantees directly. In between these extremes, the basic provider can deliver various combinations of price and/or service guarantees. To describe the scenario, we produce a specific instantiation of the generic “risk broker” of the previous chapter, making explicit assumptions on interfaces between stakeholders.

One single stakeholder might very well undertake both basic transport and guaranteed service provision. (Indeed the proposed separation has never been seen before!) Creating two separate economic entities will however be convenient for explicit economical modelling of the provision of service guarantees. Also, basing the underlying transport service on dynamic pricing makes such a separation possible and relevant. The ideas leading up to a specification of the “guaranteed stream provider” come from the work of Gibbens and Kelly [8], [7].

5.1.1 Risk broker instantiated

As stated above, the guaranteed stream provider (GSP) will be developed from the more general “risk broker” of the previous chapter. This task will be completed by considering all interfaces, and defining the interactions taking place over them to a convenient level.

One observation on the GSP might be relevant before proceeding with the requirements specifications: The refinements of the service taking place in the GSP resembles the ones taking place in a “layer” in OSI terminology. From the previous discussion, we had to assume that information stream went through the GSP on its way from the ISP to the customer. Inside the GSP domain there will probably be peer to peer interactions between entities residing close to the communication endpoints where the customers involved in the communication access their service. The GSP will thus be represented as a “layer” in the following figure, which will be used as a reference.

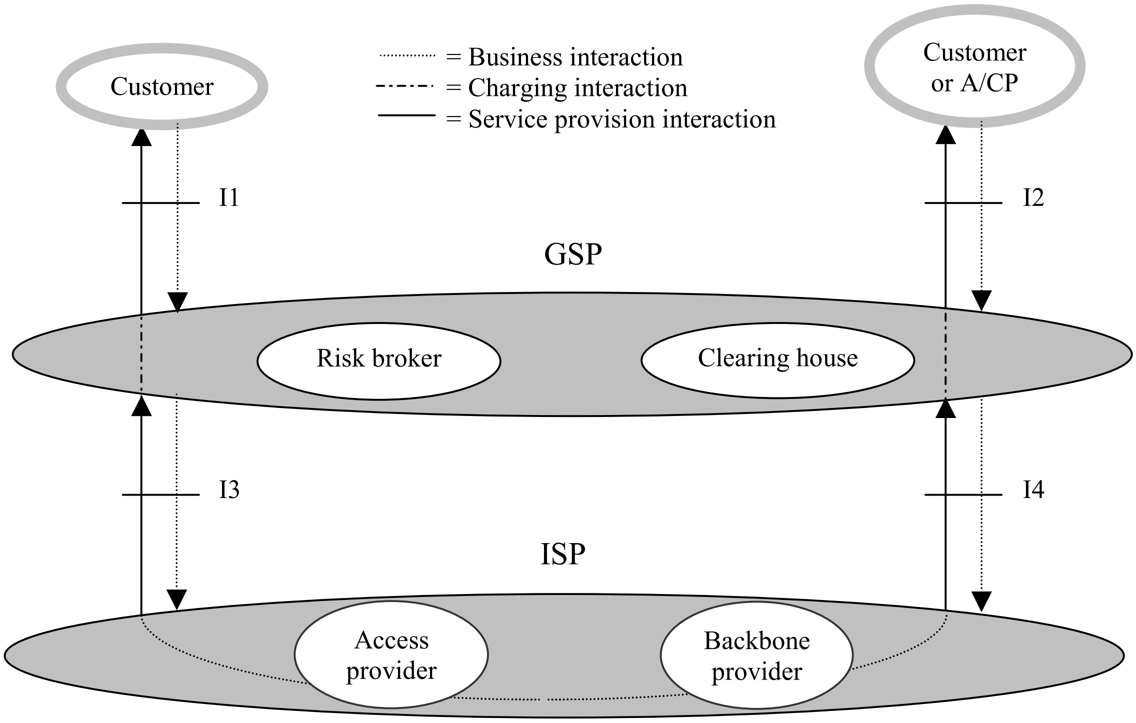

Figure 5.1 Guaranteed stream provider (GSP) scenario

The role model is based on the following assumptions. There is no switched infrastructure. Instead, the basic communication production platform is based on the packed-switched IP technology. We call the stakeholders running these platforms for ISPs although they will probably be traditional operators with a renewed production platform. Access provision is handled by these ISPs, using cable network or xDSL/fiber. They will also retain some form of usage based charging on the basic communication services they offer. The access networks are terminated in well-defined and fairly strictly regulated broadband access points (BAP).

The guaranteed stream provider also runs a charging system in order to charge users for the services offered to them. We make no assumptions here on the architecture of the GSP. It may offer its services at the BAP, using traditional signalling from the user to itself through the access network, or it may have a fully distributed architecture where services are activated at a service access point inside customer equipment as described in [2]. Other architectures may also be relevant.

5.1.2 Requirements, option a

Interfaces I1/I2, service provisioning

The services over the interfaces are unidirectional point to point IP variable-length-packet streams. The streams are explicitly set up and torn down so that they have a well-defined duration. The specific manner in which streams are controlled is a matter of architectural design. With each stream is associated a traffic descriptor, as specified in [15]. In addition to a traffic descriptor, customers will give a maximum duration for the stream requested. For traffic conforming to the traffic descriptor and within the agreed maximum duration, GSP guarantees on performance are given. The general nature of the guarantees will be that packets are delivered between endpoints within a given time limit with a given probability. More elaborate constructions are for further study. Which specific guarantees that can be fulfilled is a matter of theoretical and experimental investigation within the M3I project. At the end of the agreed maximum duration period, the service is renegotiable. It is assumed that this maximum period is substantially longer than an average call duration, and that neither the provider nor the customer will renegotiate the service during a call within the agreed time limit.

Interfaces I1/I2, charging

No charging interactions take place between the GSP and Customers or Application/Content providers. If charging functionality is distributed to execute in their equipment as described in [2], it will still be considered that the charging function is under GSP control and resides in the GSP logical domain. A scenario variation can be envisaged where customers control the charging of GSP services, and offer to transfer charging information to the ISP on request or after some other agreed fashion. This variation will not be further pursued within this specific scenario.

Interfaces I1/I2 business interactions

Customers are given all relevant information on tariffs electronically in a suitable format. The transfer of this information can be on request or provided by the GSP on a regular basis. During active stream connections, customers will be given periodically updated (e.g. every five seconds) advice of charge for the stream.

Suitable mechanisms will exist that allow customers to negotiate the divisioning of payment for the stream. This negotiation can in addition to the communicating parties also involve a third party that has a customer-relationship with the GSP.

The tariffs for the guaranteed stream service will be of the general form $C = aT + bV + c$, where T=time and V=information volume. The coefficients a and b will be derived from the traffic descriptor as described [10]. Tariffs are normally fixed, but can be revised at some infrequent interval.

The actual means of monetary transfers over the interface is outside the scope of this specification.

Interfaces I3/I4 service provisioning

The ISP offer to the GSP a best-effort datagram service. All packets delivered over the interfaces are given the same priority within the ISP's network. The ISP will however implement an ECN marking scheme. The specific ECN marking scheme used is a matter of investigation within M3I. One possibility is to use the scheme described in [8].

Interfaces I3/I4, charging

No charging interaction will take place between the ISP and the GSP. The ISP will use its own charging function to collect usage data.

Interfaces I3/I4 business interactions

GSPs are given all relevant information on tariffs electronically in a suitable format. The transfer of this information can be on request or provided by the ISP on a regular basis. GSP will be given running advice of charge for packet transfers over the interfaces.

The tariff to be used for communication is a charge per ECN-mark received. Charges apply to the receiver of packets. If more than one ISP is involved, the bi-lateral accounting is done on the basis of the difference in the number of received ECN marks.

Tariffs do not change often (e.g. once per year, on the time-scale of changes in network configurations).

5.1.3 Requirements, option b

Everything as above except:

Interfaces I3/I4 service provisioning

The service offered over the interfaces are prioritised transfer of datagrams. The priorities are set by the sender of packets according to [16]. The sender is free to choose priority level on an individual packet basis. The ISP guarantees only that packets with a higher priority level are given a higher priority in transfers than the packets with lower priority level. The number of priority levels and the specific priority mechanism to be used is for further investigation within the project.

Interfaces I3/I4 business interactions

The tariff to be used is a charge per volume in each priority class. The charge per volume increases in increasing priority, but the specific relation between priorities and priority charges is a matter of further investigation within the project. Any combination of sender and/or receiver pays should be supported.

5.1.4 Requirements, option c

Everything as in section 5.1.2 except:

Interfaces I3/I4 service provisioning

The service offered over the interfaces is a guaranteed transfer of datagrams. For a negotiated traffic specification, the ISP guarantees certain per-flow performance parameters, e.g. a guaranteed service rate and delay. The service model may be similar to the ones defined in the Integrated Services context [17].

Interfaces I3/I4 business interactions

The tariff to be used will have a dynamic aspect, but its specific nature is a matter of further investigation within the project. One particular choice is an auction-based tariff, which effectively implements price-based load control for a Guaranteed service [18].

5.1.5 Requirements, option d

Everything as in section 5.1.4 except:

Interfaces I3/I4 business interactions

The tariff will be of the general form C = f(Q,T), where Q represents the amount of resources which are effectively set aside to deliver the service and T=time. This would mean that service and price are both fixed and guaranteed before the session starts. This would certainly mean that the ISP is effectively the GSP - the traditional business model with no economically independent entity as GSP. The value of this sub-scenario lies in being a good absolute reference point for comparing the more progressive business models implied by either a varying service quality or fluctuating prices instead of solely relative comparisons between these.

5.1.6 Requirements, option e

The motivation for this variant is to cover what is expected to become a common way of realising guaranteed services in an Intserv/Diffserv environment.

Everything as in section 5.1.2 except:

Interfaces I3/I4 service provisioning

The service offered is a single Diffserv SLA (see [19]). SLAs can however be dynamically renegotiated, e.g. in terms of capacities between end points.

Interfaces I3/I4 business interactions

The service can be either fixed-priced or spot-priced. In the latter case, SLAs must be associated with maximum durations in order to avoid hogging of resources obtained during low-load/low-priced periods.

5.2 Direct dynamic pricing

While the guaranteed service provider of the previous section hides the dynamic pricing aspect of the transport service from end customers, a more obvious alternative is to offer the dynamically priced service directly to them. End customers will then need the instrument of a “dynamic price handler” as introduced and described in section 4.2. Using this instrument, end customers will have the possibility of tailoring the deployment of the dynamically priced transport service to their specific needs. The core functionality of a dynamic price handler is thus to control user’s network service class choices, according to a dynamically varying corresponding price and the utility function of the user (expressed in some degree through a specific policy).

5.2.1 Dynamic price handler instantiated

The functionality of a dynamic price handler can be divided into two separate levels (see Figure 5.2):

1. at a low level (reaction in a fast timescale), the adaptation of the service class (e.g. traffic rate, diffserv type-of-service, etc.) function, namely **QoS Control**, according to the varying network conditions (price, QoS received).
2. at a higher level (reaction in a slow timescale), the guidance of the service adaptation function to achieve an objective, which can be expressed by a utility function, namely **Price Reactor**.

The QoS Control function is responsible to realize the Price Reactor's decisions, expressed through a specific willingness to pay. It adapts the traffic characteristics (rate, priority class, etc) according to the varying network prices, in order for the charge per service unit to be equal with the corresponding willingness to pay. Its functionality could be fully implemented at the sender domain or could be distributed between the two ends as discussed in section 4.2. The part of the QoS Control function that is responsible to actually "execute" the adaptation of the traffic according to the decisions of the Price Reactor's is however always located at the sender side. Nevertheless, this adaptation could be driven by the receiver's QoS Control function through a specific protocol. In the direct dynamic pricing scenarios, we will make the simplifying assumption that the QoS Control function is fully implemented at the sender side.

The Price Reactor can be seen as the user's "intelligent agent", which is responsible to apply the user's policy according to the varying network prices or QoS. It monitors the QoS received, and sets the willingness to pay parameter of the QoS Control function based on the current price and user's utility function. Additionally, the Price Reactor's decision algorithm could also be based on the specific characteristics of the corresponding application and their relationship to the user's perception of QoS. For example, a Price Reactor could take into account the assumption that users prefer to have a single long degradation of video quality rather than many short ones.

The functionality of both the Price Reactor and QoS Control is dependent on the underlying network services. We consider three network service options offered, as in the Guaranteed Stream Provider scenario (section 5.1):

a) Best-effort datagram service with ECN marking
b) Prioritized transfer of datagrams with charging per volume in each priority class
c) Guaranteed transfer of datagrams with auction-based tariffs

According to the specific option, the dynamic prices are either explicitly set (e.g. ECN marks) or are implied from the varying QoS received for a fixed price. We will describe the different requirements that each option poses to the presented role model in the following sections.

In Figure 5.2, the corresponding role model is presented in the most general case where both parties are paying for the communication. This is may not be the most probable one, especially when the sender is an application provider, but it is the richest one, concerning the roles and interactions involved. Some interfaces have different semantics (or may even don't exist) according to the assumptions made about the distribution of the functionality of the different modules, and the paying party, as explained later in this document (see Section 5.2.5).

As far as the ISP characteristics and the requirements of interfaces I5 and I6, are concerned, the role model is based on the same assumptions (for each option) that are made in the GSP scenario (interfaces I3 and I4 respectively – see section 5.1). The distribution of charges, where necessary, is assumed made through a clearinghouse. We do not consider the various options that may apply for the bundling of application and transport services.

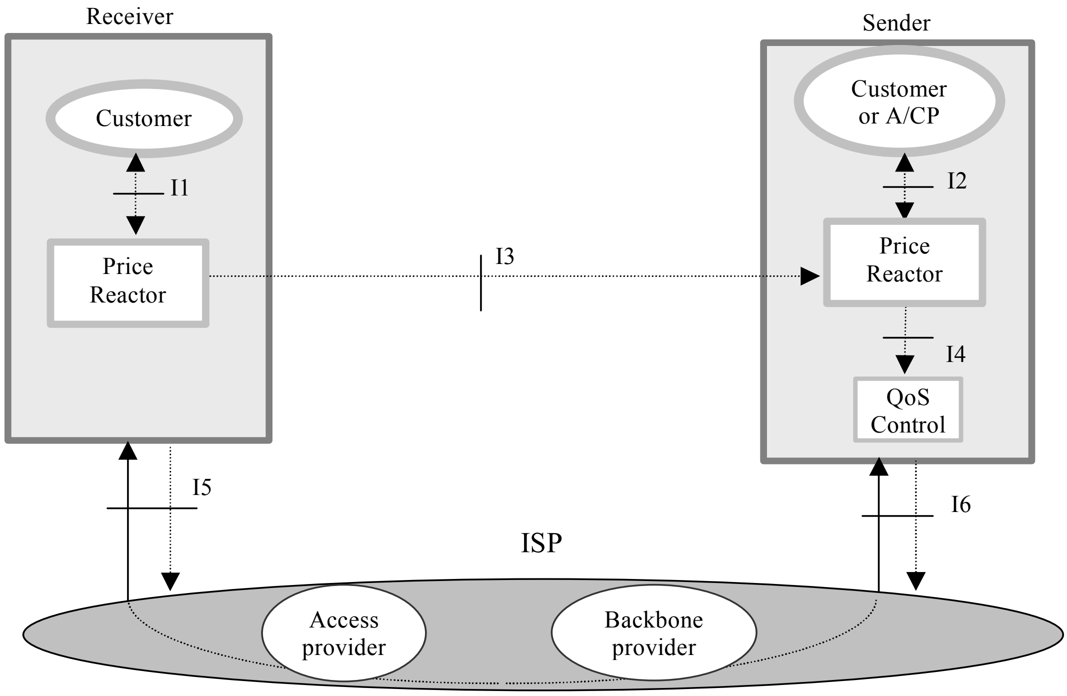

Figure 5.2 Direct dynamic pricing scenario

5.2.2 Option a

In the case of best-effort service, the network charges customers at the receiving end by the total number of ECN marks they receive multiplied by a fixed price per mark. The QoS Control module at the sender side adapts the traffic rate of the flow so that the flow's charging rate does not exceed the corresponding user's willingness to pay (one could see it as an enhanced TCP algorithm). The Price Reactor module monitors the QoS level received and the corresponding prices, and sets the calculated willingness to pay to the QoS Control.

Interfaces I1/I2, business interactions

The user expresses its utility function through a specific policy to the Price Reactor, and the Price Reactor informs him of his current (or final) charge throughout a session.

Interface I3, business interactions

The two Price Reactors negotiate in order to agree on a common willingness to pay. For example, the user who has more value could cover the extra cost of the other.

Interface I4, business interactions

The Price Reactor sets the willingness to pay parameter to the corresponding QoS Control function. That is to say the amount of money according to which QoS Control must adapt the sending rate in relation with the ECN marks received.

5.2.3 Option b

In the case of prioritized transfer of datagrams according to a finite number of QoS classes, the prices per volume in each priority class are fixed, but the QoS received in each priority class varies dynamically.

The QoS Control functionality depends on the nature of users' utility functions. If these are expressed as a given willingness-to-pay rate per given information rate, the combined task of the components of the dynamic price handler will be to match attainable charge/information rates as closely as possible to this. Other possibilities might exist as well, and the specific nature of Price Reactor and QoS control functions will need further attention in the project.

The semantics of interfaces I1 to I4 are the same as in option a.

5.2.4 Option c

Everything as in Section 5.2.2 except:

Interface I4, business interactions

The Price Reactor adapts the negotiation service according to the current price and the utility function specified by the user.

5.2.5 Scenario variations

The basic scenario presented in Figure 5.2 assumes that both parties are paying for the communication. This case is possible when both sender and receiver are end-customers. The most probable case is however that only one of the two parties will be responsible to pay. The role model could be simplified with regard to the party that is paying the bill, as explained below.

In order to achieve the situation of sender or receiver pays for the communication, it might be necessary to perform Charge Function Interactions between them.

Receiver pays

This scenario is well suited when an end-customer (receiver) is asking for a service from an application provider (sender) and he is responsible to pay for the communication costs to the Application Provider's ISP.

In this case, the receiver's Price Reactor is responsible only for the receiver's policy and the Price Reactor module at the sender side is not needed (see Figure 5.3). Nevertheless, the receiver's Price Reactor functionality could be distributed between the two ends. The knowledge of the user's utility would be still located at the user's Price Reactor. At the other side, the Application Provider can offer some basic policies (handled by his own Price Reactor) that the user's Price Reactor can synthesise in order to achieve a particular objective (see Figure 5.4).

In this case, the Application provider could bundle the application charge with the receiver's network charges with the help of a clearing house.

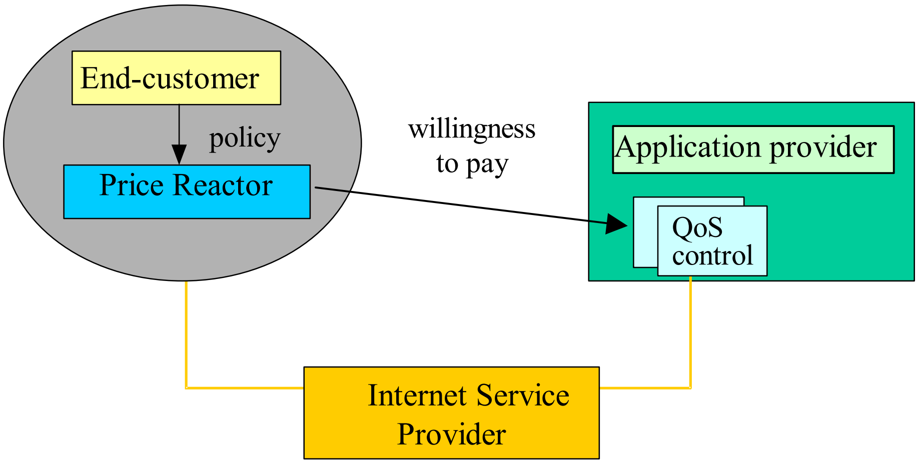

Figure 5.3 Direct dynamic pricing - Receiver Pays (1)

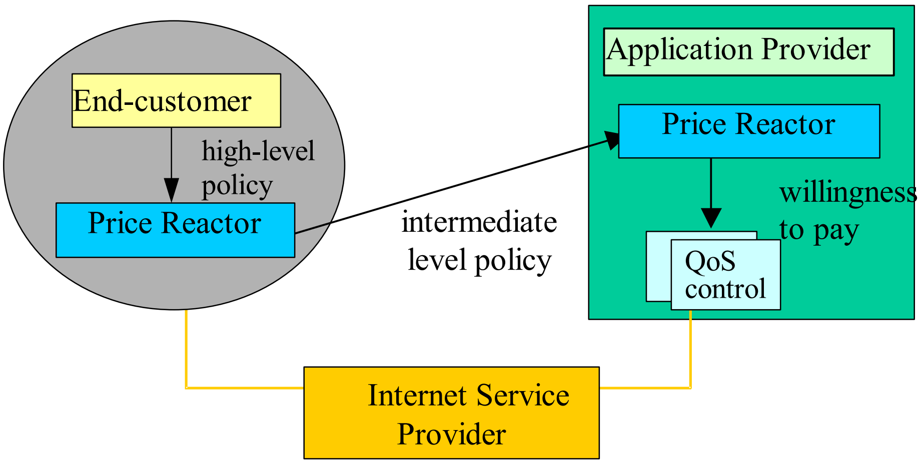

Figure 5.4 Direct dynamic pricing - Receiver Pays (2)

Sender pays

This scenario is best suited for the case where the sender is an Application Provider, which is responsible for the communication costs for both ends. This seems to be a scenario that will be immediately relevant to the industry. The willingness to pay of the Application Provider is fairly easy to calculate as it depends purely on the margins it is making from the products it bundles with QoS. If the cost of QoS reduces its margin for each product below its own threshold for that product, it stops offering that product. We could improve the policy by including the ability to discriminate between customers, applying admission control less freely to new customers, or loyal customers, etc. In this case the full functionality of the DPH is implemented at the sender side (see Figure 5.5).

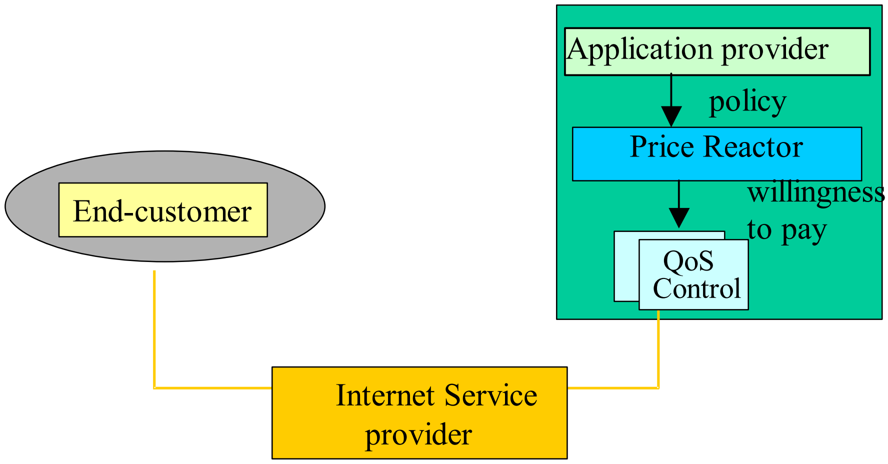

Figure 5.5 DPH Direct dynamic pricing - Sender Pays

6 General requirements revisited

A main potential in the M3I-project lies in the combination of network market mechanisms and an open architecture for systems of cooperating stakeholders. These two aspects complement each other, as a market mechanism on one hand depends on effective feedback control loops that can be realised by an open architecture, and employing such an open architecture on the other hand will contribute as a factor in the formation of competitive markets for Internet services. The specific market mechanisms that will be considered in the project are specified in the scenarios of the previous chapter. The open architecture is a consequence of the requirement that a role model as described in section 3.1.1 will be supported, where stakeholders not only cooperate in service provisioning, but are able to perform computerised interactions with regard to charging and business policies as well.

The objective of this chapter is, with the above in mind, to pull together the general requirements posed, the role model with its various components and functions, and the specific scenarios constructed with the use of the role model. The role model is of course a fiction, and can as such not satisfy any real-life requirements in itself. The purpose of relating the role model and scenarios to the requirements are therefore rather to point out some consequences of posing the requirements, and constructions and concepts that may be guides towards their fulfilment. The discussion will also serve as a motivation for the specific scenarios that were selected for further investigation in the project.

This further work will thus proceed using the scenarios as outset for studies in several disciplinary tracks. The role model will be covered in an architecture model where scenarios described at a high abstraction level here will be specified as use cases of the architecture [20]. Selected scenario variations will then be specified and prototyped. Some of these will again be taken further and used in trials and experiments. Parallel with this development there will be performed focused user trials using selected service components, and theoretical analyses with regard to market, cost and performance modelling.

6.1 Scenario motivations

The two main scenarios described in chapter 5 highlights two important possible uses of a marked managed Internet.

In the Guaranteed Stream Provider (GSP) scenario (section 5.1) the market mechanisms are implemented in the production platform, hidden from users. It might be (and it is our hope) that a production platform based on dynamic prices for resource contention resolution would have favourable properties with regards to implementation complexity and operation & maintenance as compared to more traditional approaches. This might particularly be the case when the service differentiation mechanisms can be removed from the network kernel and put either directly in end-users equipment or at the very edge of the network only. The GSP-scenarios a. (ECN-marking, 5.1.2) and b. (priority queuing, 5.1.3) are created to investigate this possibility. Even using more traditional approaches for service differentiation in the network, it might still be that having the possibilities of dynamically setting the prices for services would give a more economically efficient production platform. The GSP scenarios c. (stream auction, 5.1.4) and e. (spot-priced Diffserv SLAs, 5.1.6) are created with this in mind. In order to draw any conclusions on the questions above, reference cases are needed. The GSP scenarios d (Intserv, 5.1.5) and e (Intserv/Diffserv, 5.1.6) serve as references to a pure Intserv approach and a combined Intserv/Diffserv approach respectively.

Whereas the GSP scenario gives a setting for comparative research, the Direct Dynamic Pricing (DDP) scenario (section 5.2) investigates the possibility of giving users direct access to dynamically priced network services. This opens the possibility of even more economic efficiency than is attainable in a GSP environment, where users after all are tied to given services. The same three dynamic pricing alternatives as above are proposed (ECN-marking, priority queuing and stream auction) in sections 5.2.2, 5.2.3 and 5.2.4 respectively.

There is no conflict arising from considering both GSP and DDP scenarios. They may be realised simultaneously in the same network, and users may be offered simultaneous access to both alternatives. It might even be possible to interwork between them as mentioned in section 4.3.1 (case 4).

6.2 Motivating requirements

It is now opportune to revisit the motivating requirements set up in section 2.1. We discuss them sequentially, even if the thematic order is not quite evident. In order to relieve the reader of the need to go repeatedly back and forth in the document, the points are repeated here in italic letters.

Requirements:

U1: *the ability to instantaneously increase quality of service (QoS) by accepting different charging rates;*
This requirement is generally covered in the DDP-scenario, where users are given direct access to the differentially and dynamically priced network services.

U2: *more effective competition in a differentiated services market;*
Competition can be seen both from the user's point of view and from the service provider's point of view. In a dynamically priced network of differentiated services, users will compete for resources, covering the latter view. The services described in the scenarios will facilitate such competition. In and environment with different stakeholder cooperating through open interfaces, chances are higher that users will experience competition between stakeholders of similar kinds.

U3: *real-time feedback and validation of charges.*
Real-time feedback of charges has been specified as a requirement of both scenarios of this document. It will be covered in the distributed architecture to be developed, and will be demonstrated in the prototypes and trial systems to be built.

P1: *the ability to change tariffs and easily communicate them to the end users within seconds;*
This ability has been specified in the scenarios as above, and will be treated likewise.

P2: *the ability to hold current QoS in the presence of bad congestion effects by communicating price changes in real-time to customers;*
This requirement is very similar to *U3* above, but implies that there will be an additional relation between charges and network states. The project will cover

this point by investigating aspects of the dynamically priced network services described with the scenarios. Investigations will be both theoretical and experimental.

P3: the ability to charge differentially for applications requiring differing QoS levels, or multicast.
The GSP scenario will create the environment of the differentially priced differentiated services as implied in this point. Specific requirements and their implications in the case of multicast will need further investigation within the project.

Research issues, whether and to what extent:

D1: the demand for Internet services, including various QoS levels can be managed effectively through a pricing mechanism.
This question will be investigated both theoretically and experimentally. In the latter respect, experiments will be made by an experimental implementation of the GSP scenario.

D2: customers can flexibly access both high and low quality services, depending on their particular application needs, instead of being limited to a single best-effort service as in the current Internet;
Investigating users' abilities to relate through their applications to differentially charged differentiated services will be an objective in the M3I user trials, as specified in [21].

D3: end users in corporate organisations can exercise similar choice, but constrained by the policy of the party that is paying.
Although the role model defines the role of an end-user network provider that is a natural executor of end-user service deployment policy enforcement, there are no specific plans to work further in that direction.

D4: ISPs can recover the costs of new services, such as voice and video, that are currently provided by different infrastructures, and hence increase social efficiency by exploiting economies of multiplexing and scale, which in turn will also provide for increased network revenue.
Differential charging for differentiated guaranteed services is an important component in cost recovery for new voice and video services. The project will investigate to what extent this can be realised in a dynamic pricing environment through the realisation of the GSP scenario.

D5: simple and scalable extensions to current technology can provide the correct incentives for the economically efficient and uncongested operation of the Internet.
This question will specifically be investigated through the a., b. and c. options of both GSP and DDP scenarios. Methods of investigation will range from prototyping and experiments to theoretical analyses.

6.3 Extended requirements

The discussion on extended requirements will be carried out on the basis of the points as listed in chapter 2.2.

6.3.1 User requirements

The user requirements consider predictability of charges, transparency and accuracy of charging, convenience, applicability and reliability of invoices and payment methods. Some of these points will be passed on “transparently” to architecture and implementation work in the project. This goes for “convenience” and “applicability”. The other requirements generally relate to the functions of clearing house, dynamic price handler and risk broker developed in chapter 4.

Considering specifically predictability of charges, the nature of such predictability depends on what users are able to predict on their own traffic generation. At least two possibilities for charge predictability can be seen: predictability given full knowledge on future traffic generation and predictability given no knowledge of future traffic generation. In the former case the tariff will be given but contain usage parameters (e.g. traffic volume or transaction duration). In the latter case, charge must be independent of usage. Both possibilities are covered in the scenarios, the former kind in the options a. – d. in the GSP scenario, the latter in the last option e. of the same.

The general function of a clearinghouse will facilitate such requirements as transparency of charges (i.e. charge components of composite services to be made visible to users) and one stop shopping (users have a business relationship with only one stakeholder even for composite services). As discussed in the presentation of the clearinghouse, architectural restrictions may however still apply. It will for example be difficult to provide charge details of each leg in a connection provided by concatenated communication service providers in an edge-pricing regime.

6.3.2 Provider Requirements

The provider requirements considered are technical feasibility, and possibilities for a variety of business models, different interconnect scenarios and bundling of ISP and application charges.

Technical feasibility can not be considered at the current abstraction level, and is a requirement that will have to be passed on to the architecture and implementation processes in the project.

The demand that a variety of business models may be supported is answered by on one hand considering the open architecture to be employed, allowing flexibility in the formation of constellations of cooperating stakeholders, and on the other hand by the specific scenarios described. These scenarios of the GSP and DDP together form a richer environment for creating businesses than the current Internet.

Interconnect and bundling have similarities with respect to charging, and the proposition is to use the instrument of a clearinghouse function in both cases. In the service domain, the components of risk brokers and dynamic price handlers can be used to create a fairly rich variety of interconnect constellations, as discussed in section 4.3.1.

6.3.3 System Requirements

System requirements include flexibility, fraud protection, legal security, and openness. As in the previous section, not all system requirements will be reflected in the role model and scenarios. This goes specifically for “fraud protection and legal security”.

Considering flexibility, the clearinghouse function in the role model is designed specifically to collect and redistribute charges for composite services between two or more paying parties. The function is specified as component in the general role definition, and may be present in any stakeholder.

As for openness, it is a basic assumption in the design of the role model that both build-time and run-time openness will be possible. This requirement, playing the part of an assumption in the role model and its scenarios, is thus passed on to the work on architecture and prototyping in the project.

6.4 Tensions between requirements

The functions and scenarios proposed will alleviate some of the tensions identified in section 2.3.

Specifically, introducing the general components of risk brokers and dynamic price handlers will relieve the tension between service quality and charge. Realising both DPH and DDP scenarios in the same network, and giving the users a choice between whether to access services directly or through a risk broker, will make possible a sound foundation for the pricing of quality. This will again provide the right economic balance between users preferring quality and users susceptible to risks.

At a more detailed level, the dynamic price handler function and the corresponding DDP scenarios will alleviate the tension between convenience and dynamic pricing. The tension can perhaps not be completely removed, but using smart algorithms as the one proposed in [22] may considerably contribute to the general acceptability of dynamic pricing approaches for users. At the same level, a risk broker will remove from end users the tension between predictability and dynamic pricing.

7 Concluding remarks and acknowledgements

The value of propositions for new business models for Internet services is extremely dependent on validity of assumptions made with regards to interrelations between stakeholders in the marketplace. The models currently found, several of which were described in this report, have been formed in an evolutionary way in the Darwinist environment of current commercial Internet service provisioning. Any new model can only be evaluated by being subjected to the same environment. The M3I project is aimed at creating new and economically efficient business models. Developing these to the point where they can be tested is however a process that takes some time and effort. A clear picture of assumptions and targets should be helpful in this process. If this report can contribute in forming such a picture on the basis of the more general principles of dynamic charging, it will be its main contribution.

The report has been produced as a combined effort of the project, where all six partners have made significant contributions. In addition to the editor, the following single persons have made specific contributions:

Telenor:	Irena Grgic, Terje Jensen, Mechthild Opperud
ETHZ:	Burkhard Stiller
HP:	Jörn Altmann, Hans Daanen, Sandy Johnstone, Huw Oliver
AEUB:	Panayotis Antoniadis, George Stamoulis
BT:	Bob Briscoe
TUD:	Martin Karsten, Jens B. Schmitt

8 References

[1] B. Briscoe, “The Direction of Value Flow in Multi-Service Connectionless Networks”, Proceedings of the International Conference in Telecommunications and Electronic Commerce, 1999.

[2] B. Briscoe, M. Rizzo, J. Tassel, K. Damianakis, “Lightweight Policing and Charging for Packet Network”, in Proc Third IEEE Conference on Open Architectures and Network Programming (OpenArch 2000), pp77-87, Tel Aviv, Israel (26-27 Mar 2000)

[3] IETF, RFC-2638, K. Nichols, V. Jacobson, and L. Zhang, “A Two-Bit Differentiated Services Architecture for the Internet”

[4] G. Fankhauser, B. Stiller, C. Vögtli, B. Plattner, “Reservation-Based Charging in an Integrated Services Network, Proceedings INFORMS, 1998.

[5] M. Karsten, J. Schmitt, L. Wolf, and R. Steinmetz, “An Embedded Charging Approach for RSVP”, International Workshop on Quality of Service (IWQoS'98)

[6] J. K. MacKie-Mason, H. R. Varian, “Princing the Internet”, in B. Kahin and J. Keller, editors, “Public access to the Internet”, Prentice Hall, Englewood Cliffs, New Jersey, 1994

[7] R. Gibbens, F. Kelly, “Resource pricing and the evolution of congestion control”, Automatica 35 (1999), 1969-1985

[8] R. Gibbens, F. P. Kelly, “Distributed connection acceptance control for a connectionless network”, Proc.16th International Teletraffic Congress, Edinburgh, Elsevier, June 1999, 941-952.

[9] A. M. Odlyzko, “Paris Metro Pricing for the Internet”, Proc. ACM Conference on Electronic Commerce (EC'99), ACM, 1999, pp. 140-147

[10] D. Songhurst (Ed), “Charging communication networks; From theory to practise”, Elsevier Science B.V., 1999, ISBN 0.444.50275.0

[11] K. Nichols, V. Jacobson, L. Zhang, “A Two-bit Differentiated Services Architecture for the Internet, IETF RFC 2638, 1999

[12] K. K. Ramakrishnan, S. Floyd, "A Proposal to add Explicit Congestion Notification (ECN) to IP", IETF RFC 2481, 1999

[13] L. W. McKnight, J. P. Bailey (Eds), “Internet Economics”, MIT Press, Cambridge, MA, 1997

[14] B. Stiller (Ed), “M3I Deliverable 4: Charging and Accounting System (CAS); Design, M3I, 2000

[15] S. Shenker, C. Partridge, R. Guerin, "Specification of Guaranteed Quality of Service", IETF RFC 2212, 1997

[16] K. Nichols, S. Blake, F. Baker, D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", IETF RFC 2474, 1998

[17] R. Braden, D. Clark, S. Shenker, "Integrated Services in the Internet; Architecture", IETF Informational RFC 1633, 1994

[18] P. Reichl, G. Fankhauser, B. Stiller, "Auction Models for Multi-Provider Internet Connections", in Proc 10. GI/ITG Fachtagung Messung, Modellierung und Bewertung von Rechen- und Kommunikationssystemen (MMB 99), Germany, 1999

[19] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss, "An Architecture for Differentiated Service", IETF Informational RFC 2475, 1998

[20] B. Briscoe (Ed), “M3I Deliverable 2, Architecture; Part I, Building Blocks, Interfaces & Compositions”, M3I, 2000

[21] D. Hands (Ed), “M3I Deliverable 15.1, Trials and Experiments Evaluation Criteria”, M3I, 2000

[22] C. Courcoubetis, G. D. Stamoulis, C.Manolakis, F. Kelly, "An Intelligent Agent for Optimizing QoS-for-Money in Priced ABR Connections", To appear in Telecommunication Systems, Special Issue on Network Economics

[23] Project P806, "P806-GI Deliverable 1, A Common Framework for QoS/Network Performance in a multi-Provider Environment", Eurescom, 1999

9 Appendix: Requirements on M3I by payment systems

In M3I we decided to factor out various payment systems and assume that the architecture will be capable of dealing with any current or future system as these systems warrant a separate project by themselves. We do however need to make sure that the requirements they pose on our system are met and that the assumptions we make are warranted.

The M3I architecture connects to an external charging/payment system through the interface that notifies it of the incurred charges and the possible feedback to M3I that the payment is made or the client is still credit worthy.

The payment systems can be categorised along three types of payment: pre-payment, pay as you go, post payment or billing. The differences in these systems depend on the relation a customer has with the payment/service provider and who runs risks. The requirements the different methods pose on the M3I architecture are mainly related to the speed of notification of the incurred charges and the speed of responding to feedback regarding payment or credit.

- "pre-payment" means that the customer deposits an amount of real money with the provider that can check the balance the risk is here mainly with the customer. The relation between the provider and customer is generally of a longer-term nature.
- "pay as you go" involves some sort of e-payment or micro payment system where the customer transfers electronic cash on the fly to the payment provider. This electronic cash is not issued by some third party trusted by both parties. As far as interaction between the M3I system and the Payment system is concerned this is probably the most demanding as the payments made are generally quite small and the speed with which the service has to be stopped once the payment stops is very high.
- "post-payment" or "billing" is the more traditional way to pay for communication services. The customer gets a regular bill for the provided services. The risk is here mainly with the provider who can cut off the services if the customer doesn't pay.

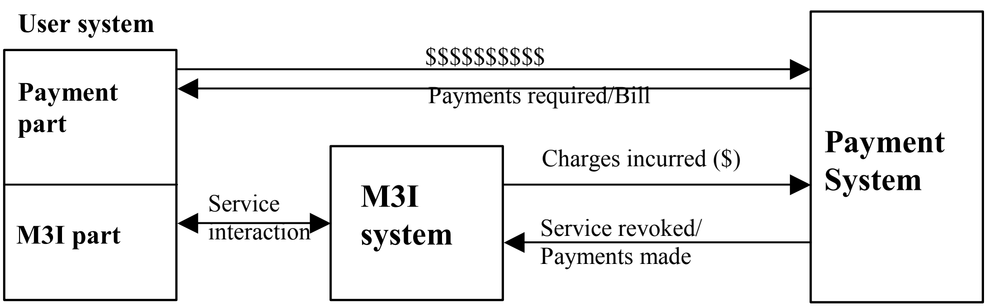

The requirements these systems pose are mainly related to the speed of notification of the incurred charges and the speed of responding to feedback regarding payment or credit.

In this picture is assumed that the M3I system does the rating and only real incurred costs are passed on to the payment system. The three payment methods give the following requirements on the performance of the M3I system.

Payment Method	**Pre paid**	**Pay as you go**	**Post paid**
"Charges Incurred" Update frequency	High to medium feedback frequency (higher as credit gets lower), low latency	High feedback frequency, low latency (this method presumably uses micro payments, so credit is always low)	Low feedback frequency, high latency (this is the traditional billing system)
"Service Revoked" response	Medium to fast (risk is shared between csp and customer as the customer prepays)	Fast response to service revoked signal, in this case it may even be necessary to provide an explicit "service allowed" or "payment made" as the trust relation between the CSP and the customer is low	Slow response as the whole feedback loop to the customer can take a long time anyway

The exact interface between the M3I system and the payment system doesn't need to be specified if the required general behaviour of the M3I system is as described above. Given the extreme high frequency of updates and the required response time to stop the service in case of non payment I think that in the case of a micro payment system the payment system can't be factored out and may have to have lower level interfaces into the M3I system. Also the Pre-paid option has to have a low latency feedback loop in order to determine when the service has to be stopped. Although in this case a larger uncertainty may be acceptable because of longer term relationships between customer and provider

An interesting architectural question that needs to be considered in more detail is who advises the user of the real cost the M3I system or the Payment system. The payment system might decide to give a customer extra discounts based on their relation with the customer (e.g. prompt payment or customer loyalty discount). If the payment system is to give charge advice the price-plan will have to be passed through them as well.

As mentioned before the main conclusion is that we will have to validate the M3I architecture against the response time and latency requirements, especially of the Pre paid and Pay as you go options.