
Market Managed Multi-service Internet

M3I

European Fifth Framework Project IST-1999-11429

Deliverable 4 Charging and Accounting System (CAS) Design

The M3I Consortium

Hewlett-Packard Ltd, Bristol UK (Coordinator)
BT Research, Ipswich GB
Eidgenössische Technische Hochschule, Zürich, CH
Darmstadt University of Technology, Darmstadt D
Telenor, Oslo N
Athens University of Economics and Business, Athens GR

© Copyright 2000, the Members of the M3I Consortium

*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: Charging and Accounting System (CAS) Design
Type: Public Deliverable
Editor: Burkhard Stiller, ETH Zürich, TIK
E-mail: stiller@tik.ee.ethz.ch

Origin: ETH Zürich, TIK
Doc ID: WP6-CAS-1.0.mif

AMENDMENT HISTORY

Version	Date	Author	Description/Comments
V 0.1	February 20, 2000	Burkhard Stiller	Document initial stage, outline, architecture
V 0.2	March 30, 2000	Burkhard Stiller	Scenario choices
V 0.3	April 30, 2000	Burkhard Stiller	Partner input inclusion (BT, HP), update layering, scenarios, and requirements
V 0.4	June 6, 2000	Burkhard Stiller	Security Considerations
V 0.5	June 9, 2000	Burkhard Stiller, Jan Gerke, Placi Flury, Martin Karsten	Interfaces, Components, CAS Architecture including layering and dimensions, Charge Communication, WP5/6 Architecture with TUD, Charge Communications
V 0.6	June 21, 2000	Burkhard Stiller, Jan Gerke, Placi Flury	Rearrangement of document structure, refinements according to Oslo Meeting, partner input
V 0.7	June 22, 2000	Burkhard Stiller, Bob Briscoe	Inclusion of mass customization section from BT
V 1.0	June 30, 2000	Burkhard Stiller, Jan Gerke, Placi Flury	Refinements of components, interfaces, example, security, finalization of deliverable
V 1.01	July 7, 2000	Jan Gerke	Corrections (mainly typing mistakes) based on comments from Panayotis Antoniadis (AUEB)

Legal Notices

The information in this document is subject to change without notice.

The Members of the M3I Consortium make no warranty of any kind with regard to this document, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. The Members of the M3I Consortium shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Table of Contents

1 Introduction	5
1.1 Charging and Accounting System (CAS)	5
1.2 Market Managed Three-tiered Model	6
1.3 Document Outline	7
2 Terminology	8
3 Networking Architecture for the CAS	10
3.1 Basic Components	10
3.2 Interplay of Components	10
3.2.1 <i>Todays Situation</i>	11
3.2.2 <i>Architecture</i>	11
3.2.3 <i>Influence of Metering</i>	11
3.2.4 <i>Further Aspects</i>	13
4 Tasks and Characteristics of the M3I CAS	13
4.1 Basic CAS Tasks	14
4.2 Identification of CAS Characteristics	15
4.3 Traffic and End-customer Considerations	16
5 CAS Design	17
5.1 Scenarios and Dimensions	17
5.1.1 <i>Location of Components</i>	18
5.1.2 <i>Replication of Components</i>	19
5.1.3 <i>Reliability of Components</i>	20
5.1.4 <i>CAS Scenarios</i>	20
5.2 Discussion of CAS Dimensions	22
5.2.1 <i>Location</i>	22
5.2.2 <i>Replication</i>	23
5.2.3 <i>Reliability</i>	24
5.3 Service Interface and QoS Layer	25
5.3.1 <i>Service Interface Definition</i>	25
5.3.2 <i>QoS Layer Definition</i>	26
5.4 Mass Customization	27
5.4.1 <i>Motivation</i>	27
5.4.2 <i>Structural Factors</i>	28
5.4.3 <i>Mass Customization Conclusions</i>	31
5.5 CAS Inter-provider Architecture	31
6 CAS Components	32
6.1 Data Gathering (Metering)	32
6.2 Mediation	33
6.3 Charging and Accounting System	34
6.3.1 <i>Accounting</i>	34
6.3.2 <i>Charging</i>	35
6.3.3 <i>Customer Support</i>	35
6.3.4 <i>User Support</i>	36
6.3.5 <i>Interfaces</i>	36
6.4 Price Calculation	37
6.5 Billing System	37

6.6 Enterprise Policy Control	38
6.7 Host/Gateway Agent	38
6.8 Service Directory	38
7 CAS Interfaces	39
7.1 General Interface Design Criteria	39
7.2 Mediation - CAS (M3I-3)	40
7.3 CAS - Price Communication (M3I-4)	41
7.4 CAS - Enterprise Policy Control (M3I-9/M3I-17)	42
7.5 CAS - QoS Component (M3I-10)	43
7.6 CAS - Host/Gateway Agent (M3I-13)	43
7.7 Feedback Interface (M3I-16)	44
7.8 Service Interface (M3I-18)	45
7.9 CAS - Billing Center Interface (Ext-2)	46
8 Service Example	46
8.1 Example Interaction between CAS and Price Calculation	48
9 CAS Security Considerations	48
9.1 Considered Security Requirements	48
9.2 CAS-oriented Assets and Threats	49
9.2.1 Consumer (Customer)	50
9.2.2 Provider	50
9.2.3 End-customer	50
9.2.4 End-user Network Provider (Customer Premises Network)	50
9.2.5 Access Provider	50
9.2.6 Backbone Provider	51
9.2.7 Data Center Provider	51
9.2.8 Market Place Provider	51
9.2.9 Communication Service Provider	51
9.2.10 Content Provider	51
9.2.11 Application Service Provider	51
9.3 Security Measures	52
9.3.1 Protection of Confidentiality	52
9.3.2 Protection of Integrity	53
9.3.3 Protection of Availability	53
9.3.4 Protection of Non-repudiation	53
9.3.5 Protection of Authenticity	53
9.3.6 Auditing	54
10 Related Work	54
10.1 Terminology	54
10.1.1 Charging	54
10.1.2 Accounting	55
10.1.3 Further Terms	56
10.2 Systems	57
10.2.1 The CATI Project	57
10.2.2 Charging Internet Services	57
10.2.3 Lightweight Policing and Charging	58
10.2.4 A Role Model for Charging	58
11 References	58
12 Abbreviations	60
13 Acknowledgements	61

1 Introduction

The objective of the Market Managed Multi-service Internet (M3I) project is to design, implement and trial a next-generation system enabling an Internet resource management through market forces, specifically by enabling differential charging for multiple levels of service [19].

Offering this capability will increase the value of Internet services to customers through a greater choice over price and quality, and reduced congestion. For the network provider, flexibility will be improved, management complexity reduced, and hence revenues for Internet service provisioning will increase. This type of price-based resource management pushes intelligence and hence complexity to the edges of the network, ensuring the same scalability and simplicity as of the current Internet.

The overall M3I work is subdivided into four main functional areas, where besides the Internet Infrastructure, Load Management including pricing and admission control mechanisms, and Applications and Middleware, the Charging and Accounting System (CAS) based on principles of metering and mediating relevant data will allow for the determination and utilization of network resource usage information on a per-customer basis, which provides usage-charging feedback to customers.

1.1 Charging and Accounting System (CAS)

The Charging and Accounting System (CAS) supports that applications can request authenticated reports of usage and charges on a per-session basis to enable bundling of network charges with those for other services. Traditional itemized usage charges will also be available, but whenever required rather than in weekly or monthly batches. The defined M3I work package 6 (Charging and Accounting System) will design, implement, and evaluate the technology needed to allocate charges for multi-service differential services in user and Internet Service Provider (ISP) market-places according to prices communicated [21] between provider and customers. The CAS developed within M3I will provide an advanced and flexible approach for, e.g., implementing static or dynamic pricing schemes.

The CAS' aim is to ensure that it is as generic as possible, with the ability to configure appropriate algorithms that are specialized to measure different aspects of the network and of services provided over this network. The framework for service-oriented charging possibilities is developed, including different levels of granularity ranging from packet level, if required, to session levels. Of particular interest are higher level services, which are provided typically in an end-to-end fashion. An important new aspect of the CAS is its potential to provide the means to setup a price-controlled feedback loop between service usage, its information to customers or users, and a return path to the network elements.

The framework includes measuring requests for network service, the determination of accounting records measuring the actual delivered service, the calculation of charging records, as well as the preparation of billing records. Security mechanisms are considered for the CAS on the design level, both to avoid fraud and also to address the important distinction between various types of these data records being utilized within M3I. The CAS itself is neither concerned with billing actions nor with settlements, but it will provide appropriate interfaces for billing data to be processed at a later stage with well-known billing systems. For the integration of electronic payment schemes a potential interface will be available, however, depending on the efficiency of a suitable technical solution, technology-dependent details need to be taken into account specifically.

1.2 Market Managed Three-tiered Model

The infrastructure of the market managed multi-service Internet is based on an overall three-tiered model (cf. Figure 1). This model outlines the basic states and sources of information within these three distinct layers as well as its interfaces to the to be designed Charging and Accounting System (CAS). This type of information and these interfaces are essential for the design of the CAS.

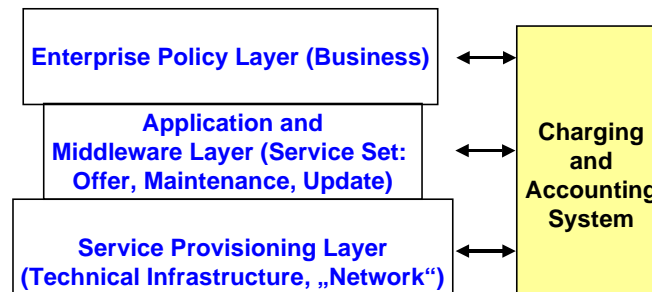


Figure 1: Market Managed Three-tiered Model

Starting from the topmost layer, customers and providers within an Internet services market interact for any type of business based on business models defined within the **Enterprise Policy Layer**, which defines, amongst others, the products to be exchanged, models of business interactions between customers and providers, pricing mechanisms, and agreements upon an offer. Details of relevance for the CAS may encompass, *e.g.*, rebate systems, discounting schemes, service plans, or service pricing models. These details form the business-dependent and business-central policy, which may not be published as such but is required to provide the CAS with operational dimensions. However, to perform any type of market-driven enterprise policy, the CAS needs to offer a fully flexible set of service descriptions, which are applicable to all areas of enterprise policies.

Besides these necessary business perspectives the technical view of the market managed approach is founded in the **Application and Middleware Layer**. It provides functions or policies, which are (1) initiated due to a pre-determined application or (2) acting on the application's behalf, such as a given enterprise policy. Within this layer, a set of (value-added) communication services is provided, utilized, and charged according to customer demand. Middleware abstracts from the details of the technical infrastructure of the network itself. The middleware is able to provide a generic service set for offering, maintaining, and updating all types of communication services. Therefore, the CAS interface includes application-centric configuration options for particular session and services descriptions as well as more generic service descriptions according to the middleware layer functionality.

Finally, the networking tasks of service provision as identified in the **Service Provisioning Layer**, form the basic interface for the lower layer of the new Charging and Accounting System. Therefore, data and information on the technical infrastructure, the "network", are collected and maintained depending on the service offered by the middleware layer.

Earlier work on charging and accounting in telecommunication systems has been focused on connection-oriented networks, such as the telephony network, Asynchronous Transfer Mode (ATM)-based networks, or leased lines. The Internet on the other hand provides a connectionless network layer and an Internet Protocol (IP)-based network service. While the set of traffic modeling parameters and service parameters for connection-oriented networks are quite well understood and agreed upon, these parameters remain heavily debated for the Internet. *E.g.*, the inter-packet arrival time for an Internet service makes a

significant difference for this service. However, how should a future system account for this parameter? In addition, as for connection-oriented networks the call blocking rate determines the level of utilization for a given topology and the potential sender is blocked in sending data into the network, connectionless networks suffer the problem of congestion, since in general there is no admission control available. Of course, a set of newly defined Internet services proposes the existence of such an admission control, however, a commonly agreed upon architecture has not been developed up to now. Once congestion situations can occur in a network, congestion control mechanisms are required. Traditionally, these mechanisms have been operating in the pure technical domain (e.g., by dropping packets), but left out incentives to evaluate the requested service by any economic measures. Sensible pricing of services in an open services market is required additionally [33]. However, this approach depends on the technical ability to collect and account for the data necessary to charge the customer. Therefore, the developed CAS will provide the technology required to support pricing-based mechanisms for congestion control.

At this point M3I and the CAS will start to supply advanced technology, which integrates technical and economic congestion control measures and mechanisms for tomorrow's Internet. This market managed approach requires the fine-grained investigation of appropriate, efficient, and flexible solutions for the Internet, since its multi-service shape is not completely defined and may change over time. Therefore, the need to provide most generic, but most efficient solutions at the same time is a key optimization dimension.

Based on these basic thoughts, the internal structure of the CAS as well as required interfaces to other components of the M3I architecture are developed, designed, and documented in this work. The CAS as well as its basic characteristics are defined as part of the requirements specification for M3I as a whole according to [1]. In addition, the required understanding of terminology, business roles, and technical components is defined in the same document and applied for the CAS developed below.

1.3 Document Outline

This document on the Charging and Accounting System (CAS) Design is organized as follows. General terminology utilized within this document and M3I is defined in Section 2. The CAS is integrated in a precise networking model for the Internet including network elements such as routers and hosts, which is presented in Section 3. Based on this discussion the set of tasks and required basic characteristics for the M3I CAS is discussed in Section 4. The CAS design is presented in Section 5. It includes CAS design impacts in terms of details for alternative scenarios, CAS component dimensions to be considered, and an inter-provider view. Afterwards, the details of the CAS components are designed in Section 6 and its interfaces are outlined in Section 7. While an example service scenario and the use of CAS components is presented in Section 8., Section 9 discusses security requirements for the CAS. Finally, Section 10 summarizes related work in terms of existing CAS approaches, their ideas, and an overview of related areas, such as terminology and systems. The list of references and abbreviations utilized in this document concludes the CAS design work.

2 Terminology

The following alphabetically ordered list of basic terminology defines the required terms to understand a concise and well-defined CAS design. It includes the definitions taken in [22].

- **Accounting:**
Summarized information (accounting records) in relation to a customer's service utilization. It is expressed in metered resource consumption, *e.g.*, for the end-system, applications, middleware, calls, or any type of connections.
- **Accounting Record:**
An accounting record includes all relevant information acquired during the accounting process. Its internal definition is for further discussion.
- **Billing:**
Collecting charging records, summarizing their charging content, and delivering a bill or invoice including an optional list of detailed charges to a user.
- **Billing Record:**
A billing record includes all relevant information acquired during the billing process. Its internal definition is for further discussion, however, not part of the M3I work.
- **Charge Calculation:**
Completing the calculation of a price for a given accounting record and its consolidation into a charging record, while mapping technical values into monetary units. Therefore, charge calculation applies a given tariff to the data accounted for.
- **Charges:**
Charges determine the amount of monetary value that needs to be paid for a particular resource utilization. It is contained in a charging record.
- **Charging:**
The overall term "charging" utilized 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. Note, that billing as such is not included in this definition.
- **Charging Record:**
A charging record includes all relevant information acquired during the charge calculation process. Its internal definition is for further discussion.
- **Costs:**
Costs determine the monetary equivalent on equipment, installation, maintenance, management, operation of networks, network entities, and service provisioning. Many different types of costs can occur but it is important to note that in the case of CAS only costs in terms of money are of interest.
- **Customer:**
The role of the customer identifies the focus of an (end-)user or an enterprise in its business behavior. This means that purely economic aspects are relevant in this case, such as a service selection based on an optimization of service utilization and cost.
- **Mediation:**
In most cases the data which is collected by metering is very technical data. *E.g.*, it contains information about packets and queue lengths but doesn't relate this data to specific customers. Mediation transforms this data into a form which can be used for storing and further processing. This is done, *e.g.*, by collecting and merging data from several metering units and other sources like a list of customers and their IP addresses.

-
- **Metering (Data Gathering):**
Determining the particular usage of resources within end-systems (hosts) or intermediate systems (routers) on a technical level, including Quality-of-Service (QoS), management, and networking parameters.
 - **Price:**
The price determines the monetary value the user owes a provider for his service provisioned and utilized, in particular it is the price per unit service. It may be based on charges and costs or it may be determined by other marketing means.
 - **Pricing (Price Setting):**
The specification and the setting of prices for goods, specifically networking resources as well as services in an open market situation. This process is part of the enterprise policy layer of the M3I architecture, but requires appropriate communication means being in place, which are provided by M3I work package 5.
 - **Quality-of-Service (QoS):**
QoS defines the quality of a service provided, it contains technical application-level as well as network-level views and definitions. Its particular specialization for M3I is for further discussion, however, a commonly agreed upon definition will be taken from other work. With this respect the definition from ITU-T, E. 800 [17] is applied initially: "The collective effect of service performance which determines the degree of satisfaction of a user of the service".
 - **Service:**
A service enfolds autonomous and network dependent tasks needed for application execution. An application typically employs several and presumably distributed services to provide full functionality¹.
 - **Session:**
A session consists of one or more (virtual) connections between hosts, which communicate with each other. It is characterized by a clearly defined starting and end time. Sessions are supported by of one or more parallel or consecutive executed services.
 - **Tariff:**
The algorithm used to determine a charge for a service usage. It is applied in the charge calculation for a given customer and service he utilizes to calculate the charges.
 - **Tariffing:**
The process of deciding upon the algorithm used to determine a tariff.
 - **User:**
The person who uses a network service by running applications.
 - **User Agent:**
A human being or an application acting on the user's behalf who utilizes network resources in a technical manner. This role does include the technical service utilization, but does not cover any business aspects, such as valuation of a service according to an optimal service/cost ratio.
 - **Utilization:**
The metered usage of resources in a re-producible fashion, quantified by defined value ranges and according units.

1. The term *service* differs from the service definition of WP 2.1 in the sense, that it is less generic and already focuses on a technical point of view (vs. a economical an architectural one).

3 Networking Architecture for the CAS

This section provides the basic assumptions on a suitable networking architecture for a CAS. While the identified components are scattered around in the network, particularly the multi-provider Internet network, their basic interactions are described.

3.1 Basic Components

Several technical components are needed for a charging and accounting system. As illustrated in Figure 2 and described in the following as well as in [22], the general scenario contains various interconnected communication service providers. Each provider has a network consisting of *routers* and *network links* between them, *accounting systems*, and a *billing system*. *Metering systems* are components inside the networks. They can be independent components or can be combined with routers. In either case, they generate accounting information (base *accounting records*) which are gathered and accumulated in accounting systems. The accounting systems in turn forward the accumulated and perhaps abstracted accounting information through a *charge calculation* function towards the billing system. The charge calculation translates the accounting information into *charging records*, hence, it maps the resource-oriented information from the accounting systems into monetary values. The billing system uses these values to prepare the bills to be sent to customers. Within the charge calculation, and perhaps the billing system as well, any discounting strategies, marketing-driven pricing schemes, or simply fixed prices can be applied.




Figure 2: Possible Location and Replication of Charging Components in Provider

3.2 Interplay of Components

Components need to interact to provide the offered functionalities to customers.

3.2.1 Today's Situation

In existing billing systems of today's providers, the setting of prices, the function of charge calculation, and the billing itself is integrated, even additionally combining the maintenance of service classes, user profiles, customer data, identities, and banking account data. Although the above mentioned steps still can be distinguished clearly, they are almost completely centralized within a single system. Future billing systems need to be able to integrate a variety of different charging records, even from different communication providers or content providers, since customer's demand is one-stop billing [28]. This strongly suggests dividing the existing monolithic billing systems into several components with clearly defined interfaces. By doing this it will become possible to exchange individual components and to integrate different components supporting different technologies without having to adapt the whole system.

Additionally, interfaces to Accounting, Metering and other components have also to be defined. The goal is to identify components and their relations to each other and to create an open and complete system structure which allows charging and accounting of different technologies from the data orientated level (Metering) up to the money orientated level (Billing).

3.2.2 Architecture

The conceptual separation between the various types of components and their interactions can be done as shown in Figure 3.

Data gets forwarded by routers. The usage of the routers is measured in a data gathering (Metering) unit. This usage data is then sent to the Charging and Accounting System and to the Price Calculation. Since the amount of usage data is usually too big to be sensibly processed (see Section 3.2.3) it must first go through Mediation where it is transformed into a form suitable for further processing. The raw usage data is also not yet associated with a specific user or customer. It is necessary to identify a customer responsible for each data unit sent. This is best already done in Mediation because then the amount of data going out of Mediation can be much smaller because usage data of one customer can be aggregated.

The Charging and Accounting System takes the measured usage data and transforms it into charges which again are the input for the Billing System. To be able to calculate charges the CAS needs the price for the service the user utilized. It gets this price from Price Calculation which calculates it using price models which can depend of the utilization of the router. So prices might get higher with rising data traffic. The prices are communicated back to the end-user so that there is a feedback loop which can help to solve congestion problems. There are several ways of communicating the current prices to the end-user which are shown in Figure 3 and are explained in detail in [19].

The presented components can be controlled by using the Enterprise Policy Control. This component provides an interface to the ISP by which it can set the various parameters of the other components and access information from the CAS, e.g., to generate statistics. All components and interfaces are explained in more detail in Section 6 and Section 7.

3.2.3 Influence of Metering

The metering systems deliver the basic information used for charging. They have to detect the resource usage by subscribers – which can be the actual usage for transport and which also can be the reserved but finally unused capacity. For reserved resources, information

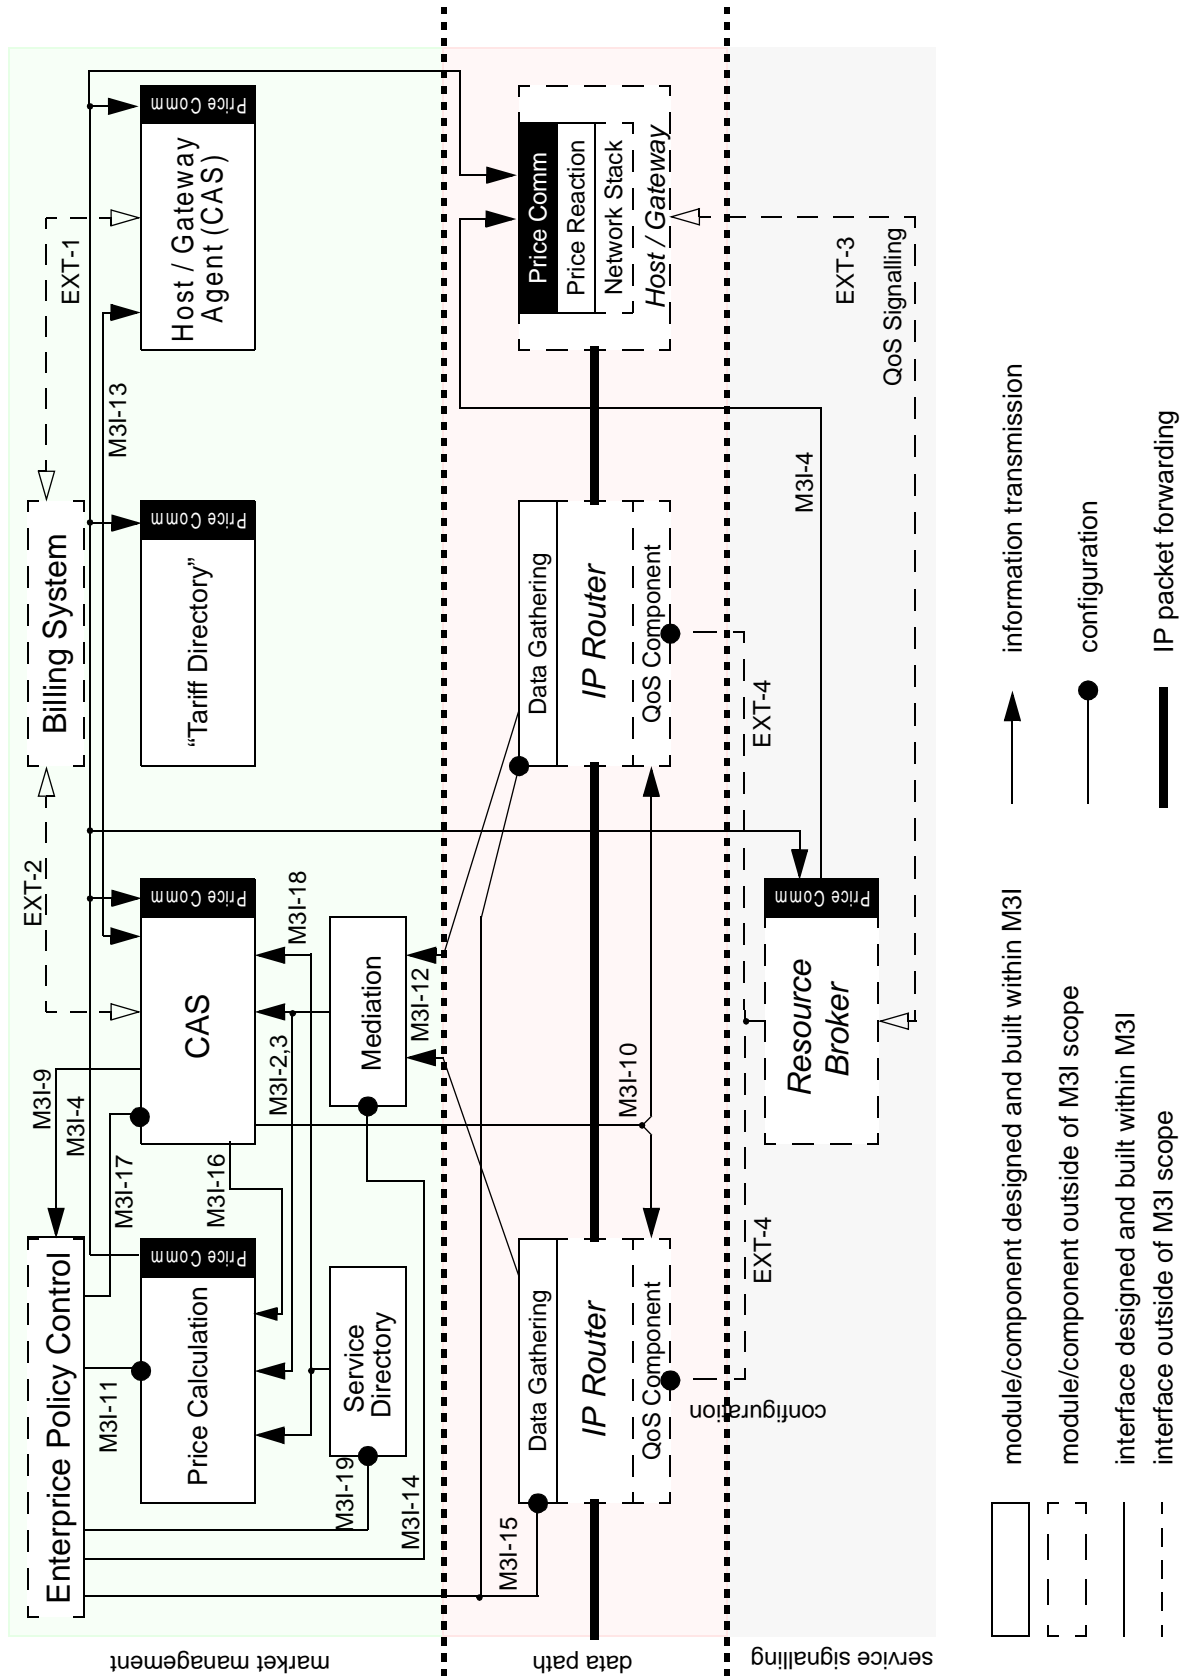


Figure 3: System Architecture of WP5 and WP6

available within routers or Bandwidth Brokers should be used. Therefore, the Protocol Data Units (PDU) exchanged among these entities may contain (policy) objects which describe announced prices and corresponding charges [8], [11], [23]. For traditional best-effort data traffic, no such explicit information about resource demand and usage is available. Therefore, and also for potential volume-based charging for reserved traffic, metering services are necessary which measure and count the transferred data volume. These units can be part of routers or additional devices attached to LANs or to metering ports of routers. Of course, depending on the performance requirements of such a router or link, using separate devices can offer further advantages besides performance in the area of reliability, division of responsibilities, etc.

Of course, the basis for a solid and acceptable pricing model is determined by various issues, such as the technical possibilities of the metering task, the performance required by this task, and its granularity of information collection. Only the data being metered may be used for the accounting record and the unit service being priced, which in turn determines the maximum billing granularity achievable. Certainly, the degree of granularity of this information may lead to a severe technical inefficiency for metering data to apply usage-sensitive pricing methods, if care has not been taken to reduce this data set to the basic essentials only.

These technical issues have been studied mainly in the context of RSVP and IntServ, so far. Corresponding architectures and components have been discussed by, e.g., [8], [23], [5]. These approaches extend the RSVP protocol to distribute charging and accounting information among routers, respectively. Additionally, protocol mechanisms to forward usage information from the metering systems to the accounting systems are required.

3.2.4 Further Aspects

By now, there are many technical aspects which have not been studied in sufficient detail in the literature. Furthermore, the interplay of the various approaches to provide quality of service makes the situation much more complicated. The discussion of all these issues is out of the scope of this report. As an example, the problems which already occur for relatively well established communication mechanisms such as IP multicast are briefly described. The treatment of IP multicast is problematic in general due to the anonymous membership model. The number and identity of participants in a particular session is unknown, yet, in case of sender charging (and also for inter-provider charging), such values influence the final bill. Hence, one provider would typically charge the other for the required effort to serve a multicast session with a specified number of participants/overall resource consumption. But, due to the anonymity of IP multicast, the effort of one provider cannot be controlled by another provider, the latter just has to trust the former.

There are a lot of further questions about security, trust and related issues which need to be taken care of as well as prerequisites for any solution. Hence, suitable security infrastructures and mechanisms must be used for the information exchange among the components.

4 Tasks and Characteristics of the M3I CAS

This section provides the overall description of basic CAS characteristics required for the M3I CAS, together with a discussion of relevant assumptions made for its design. It focuses on and exploits the context in which such a CAS resides.

Charging is the crucial feature in telecommunication services [30]. This has been determined with respect to the definition of suitable pricing schemes for a range of offered services. However, once the service is provided in a connectionless fashion, such as within the Internet, the pure economic view will not be the core focus of the problem alone.

To be able to allow for an efficient provision of multiple services within a single packet-based network, problems of appropriate accounting schemes, the difficulty of suitable parameter metering, and the open issue of service differentiation on a technical level are added. Therefore, the separation of tasks as proposed for the M3I project in “Pricing Mechanisms” in WP5 and “Charging and Accounting” in WP6 reflects the sort of economic versus technical point of view and separation of concerns. However, and that determines the most difficult and closely related features of pricing mechanisms as well as charging and accounting, these two areas are not independent at all. On one hand, this is due to the feedback loop required to charge services and service provisioning based on technical measurements and by technical means. On the other hand, this is due to the market-driven feedback which is based on technically metered data, which are collected and mapped onto different monetary equivalents.

4.1 Basic CAS Tasks

The CAS is supposed to support the following list of tasks:

- Perform service charging. Services include the ones provided by a variety of service providers (cf. [1] for their different definitions and distinctions), which are offered in an open market situation. This charging task needs to be as far as possible service-independent, to ensure future extensions and adaptations to yet unknown services.
- Perform accounting tasks according to service definitions. Data gathered from the physical infrastructure and mediated due to some policies needs to be accounted for. This requires the knowledge of “sessions”, “durations”, or “flows”. Mainly, these information are derived from the metered data as well, such as “begin-of-session” or “end-of-flow”. If such starting and end points can not be determined explicitly, heuristics need to be applied for session or flow detection purposes. In any case, the “length” of a communication relation will be recorded, if any usage-based charging approaches are to be supported.
- Perform multi-service accounting. The accounting task for a single service, which is well known is performed by an algorithm, which utilizes a clear service specification. In case of multi-service provisioning, these service specifications must exist concurrently and need to be maintained that way. Therefore, the separation of incoming data and their mapping onto the particular service in operation is essential.
- Support transport, service, and content charging. The optimal design for a Charging and Accounting System includes a combined approach for the three different levels of charging. Transport charging, sometimes termed network charging or network access charging as well, forms the basis for providing a system to deal with the transfer of data, mainly based on a general network infrastructure, such as the Internet. The service charging located on top of this level allows for the clear distinction of different services including different QoS requirements and resource consumptions. Certainly, the transport charging will be integrated into this concept and may even be completely hidden. Finally, the content charging includes the accounting tasks for information which is specifically monetary-sensitive and needs to be paid for by reading, using, or copying it. Based on the level of business interactions, it might be useful to apply content

charges for certain services only, integrating invisibly by the customers the underlying transport and services charging. Therefore, the CAS needs to be designed internally in an open fashion for multiple accounting and charging levels.

- Support different levels of security for charging and accounting information. For every data and information which is related to monetary equivalents, these data show a certain degree of sensitivity. However, based on the dedicated level of interest a single accounting record, a single metered routing data, or a charging record may not be a security problem, since their lifetime and validity, and therefore asset, are short. But other combinations of aggregated data, say flow-related information in terms of usage information, duration, and customer identification, form a critical information. The details of security and its basic mechanisms to be applied are discussed in Section 9.
- Support auditing. Communication services offered in a market environment need mechanisms which support the proof of service delivery under well-defined circumstances. Therefore, an auditing functionality will be based on accounted for data, which may specifically restricted, structured, or stored depending on legal aspects, such as telecommunications acts. While these mechanisms will not be part of the M3I work, the interface for them can be provided in a most general fashion.

These tasks can not be evaluated at this point in a very fine-grained manner, since the fine design of the CAS will identify those interfaces and protocols as well as data structures, which are to be taken care of for an efficient implementation of the CAS. However, a set of directly dependent characteristics are listed in the following Subsection 4.2. Once these tasks and basic characteristics are identified a discussion of meaningful combinations and suitable solutions is performed afterwards.

4.2 Identification of CAS Characteristics

The CAS is supposed to support the following characteristics. They include customer, provider, and system characteristics. This list is based amongst others and further extensions on [9], [23], [26], [28], [30], [32] as well as M3I internal discussions.

- Technical feasibility for Internet environments.
- Efficiency in terms of performance aspects.
- Scalability in terms of multiple services, multiple users and customers, multiple locations, and multiple providers.
- Be technology independent.
- Be soft real-time capable.
- Support multi-service networks.
- Be single-user (single-customer) oriented.
- Allow the implementation of different pricing models.
- Allow for different ISP cost models.
- Allow explicit cost recovery.
- Support the predictability of charges, where applicable.
- Support the transparency of charges, where applicable.
- Allow the support of a variety of different business models, such as service bundling, service offerings, and targeting particular (customer) market segments.

-
- Allow a full transparency of charge calculation and billing. The technical parameters need to be maintained unambiguously as well.
 - Allow interfacing to existing billing systems.
 - Accuracy of charging to be able to follow the interdependence between a particular pricing model and the accounted for resources.
 - Full flexibility in terms of flow of value, flow of information, and flow of stimulus.
 - Fraud protection for providers and customers.
 - Support of auditing mechanisms.
 - Legal security in terms of fulfilling relevant (data communication) acts and amendments.
 - Allow refunds.
 - Allow discounts.
 - Allow reporting of usage before charges are applied.
 - Allow rapid reallocation of addresses to customer identity.
 - Support the flexible service provisioning of providers.

These characteristics may be contradictory in certain scenarios. A brief discussion of suitable solutions and possible combinations will be included in the following fine design of the CAS, where data structures, protocol data units, protocols, interfaces, and detailed components are known and specified. The important reason for having these tasks and characteristics enlisted in advance, is the fact to sharpen the designer's view on potential problems and areas of crucial performance measures. Therefore, the following design of the CAS architecture, its embedding into the WP5 and WP6 areas of work, and its influences on the modeling tasks in WP4, in particular the ISP cost model will consider these tasks and characteristics as essential.

4.3 Traffic and End-customer Considerations

Besides these tasks and characteristics, the type of traffic as well as the end-customers are to be considered in more detail. This is essential for the design and outline of all CAS components with respect to their efficiency, scalability, and robustness.

Basically, at least four different types of end-customer traffic need to be considered for the CAS design. This traffic may be classified initially into the following four classes, but may be refined at a later stage for more fine-grained investigations:

- Best-effort traffic, such as e-mail.
- Guaranteed traffic with small reliability requirements and low bandwidth demands, such as audio conferencing, IP telephony, or web traffic.
- Guaranteed traffic with small reliability requirements and high bandwidth demands, such as video conferencing or Video-on-demand services.
- Guaranteed traffic with high reliability requirements, such as remote control or sensing and types of electronic-commerce interactions.

Depending on these four classes of traffic provided in the network and the characteristics specification for the CAS, the basis is determined for evaluations of functionality and performance.

5 CAS Design

The design of the CAS requires as a first step the definition of the CAS architecture which has been introduced in Section 3. Based on this definition, components and interfaces can be derived. However, to design a more detailed CAS further steps have to be taken. Of special importance is the consideration of various influences which can have a big impact on the design. In this section these influences will be discussed. This includes the presentation and discussion of different dimensions in which a CAS may vary. The close relation between QoS and market forces is also taken care of by the introduction of service interfaces and the QoS layer. Finally, the individual requirements of customers and the interconnection of providers are considered.

5.1 Scenarios and Dimensions

While the charging components for the overall CAS design (cf. Figure 2 and Figure 3) have been identified it must still be determined how these components are implemented and deployed in each possible scenario. There are a lot of potential scenarios which may include several different ISPs. There are three dimensions in which a CAS can vary according to the scenario (cf. Figure 4).




Figure 4: CAS Scenario Dimensions

When deploying a CAS for a specific ISP the ISP type defines a set of different choices based on the distinction of roles for Access ISPs and Transit ISPs according to the requirements [1]. Depending on the ISP type the location as well as replication of components will determine suitable and less useful combinations of components. However, there is no general set of criteria available at the moment depicting the optimal location and replication of components for a given scenario. However, the work on cost modeling and its aspects may determine a suitable design process for this aspect [34].

The dimension of **location** defines where the components are located. In particular, the “in-sourced” location refers to the fact that the ISP itself hosts this component and provides the according functionality internally. The “out-sourced” location defines that this component and functionality are being performed outside the scope and administrative domain of the ISP. Mainly business case assumptions and the size of the ISP considered will determine the final location of components in a given ISP infrastructure. In addition, security-relevant questions may arise, once the out-sourcing of financial activities is intended.

The dimension **replication** defines how many of the components considered exist in a given environment. Mainly the number of clients served by an ISP or the number of interconnection points with peering ISPs will determine the number of replicated components required. However, besides the pure replication an important issue is the interaction between these replicated components. Appropriate protocols (open, ISP-specific, or vendor-specific) need to be selected for a suitable and correct design and implementation.

The dimension **reliability** defines how reliable the components have to be. The needed degree of reliability depends only indirectly on the ISP Type. It rather depends on the before mentioned other dimensions of location and replication. It also depends heavily on the type of component. Nevertheless, the needed reliability of components is a dimension in which a specific CAS can differ from others and which will vary from one ISP to another.

5.1.1 Location of Components

Basically three different location alternatives exist for A, B, and C components. They are outlined in its basic principles in Figure 5.




Figure 5: CAS Scenario - Location Alternatives

Case (a) depicts the situation where all components, irrespective of their functionality are “in-sourced” within the ISP’s domain. While metering M is always performed at the particular router of interest, the accounting components A are operated and set-up within the ISP’s domain as well as the charge calculation component C. All of them do provide appropriate interfaces for service-oriented charging. However, these interfaces do not need to be publicly available, since they are operated within the ISP’s domain. Finally, the billing component B is still within the ISP’s domain, but it needs to provide open interfaces for a real business case, since there will be most likely customers residing out of the scope of the ISP’s domain which require billing data to be delivered to. Therefore, even though the billing is in-sourced, it provides the interface to the outside world of other ISPs. Case (a) reflects the most extreme situation of a complete in-sourcing.

Case (b) shows the opposite case. Certainly, in any case metering will remain with the equipment to be metered. All other components (A, B, and C) are out-sourced. Outsourcing these components may be useful for smaller ISPs, where the effort to set-up and maintain them is too costly. However, only a clear calculation of costs involved in this situation will reveal which solution is the best. In addition, since the accounting component A still reflects the technical-oriented functionality, its outsourcing may become difficult in a real situation. *E.g.*, the utilization of a network management tool with integrated accounting

functionality would be operated within the ISP's domain. Therefore, case (b) reflects the most extreme situation of a complete out-sourcing.

This leads directly to case (c), where a number of realistic assumptions on the size and the operational effort of an ISP are taken into account. On one hand, the ISP has to perform an in-sourced network management and accounting system, since the number of end-customers being served and the topological distances of its network are too large. On the other hand, the ISP does not want to perform all the charge calculation and billing functions, since it has not reached the critical mass of end-customers which would efficiently allow for running these quite complex systems locally. In the traditional telecommunications industry case, this still holds true for smaller telephone providers who utilize the old monopolist's billing and charge calculation systems.

Considering the communication channels between these cases and components, they show a commonness with respect to their interfaces. This commonness is the openness of these interfaces as soon as the interaction between the ISP's internal domain and an external partner is necessary. Due to the fact that the components designed for the CAS may be located in different places, all interfaces are required to be as open as possible, and the architecture of the overall system should not prevent this openness.

5.1.2 Replication of Components

Basically two different extreme replication alternatives exist for A, B, and C components. They are outlined in its basic principles in Figure 6.




Figure 6: CAS Scenario - Replication Alternatives

Case (d) identifies the situation with the maximum level of replication for all components within a single ISP's domain. This means that metering, accounting, and charge calculation are performed per router in a given network infrastructure. There are certainly some reasons to provide this degree of replication, especially in situations where the ISP requires a fully reliable and, therefore, redundant infrastructure. A single failure in one of these components will result in the loss of service for the set of directly interconnected components and router only.

As an extreme case on the other side of the spectrum, case (e) identifies the minimalistic solution for replication - not a single component requires any replication at all. Again, this determines the main advantage of the architecture and design of the CAS that the single occurrence of a component is sufficient to perform all the requested for and required functionalities. Of course, the questions on performance and suitability in a given case of

further topology, end-customer, and demand data can be answered only with respect to an integrated investigation of the location dimension.

5.1.3 Reliability of Components

An important question when installing CAS components is how reliable these components have to be. In most cases high reliability also means high effort which directly translates to high costs. In a market driven environment each provider must keep his costs as low as possible while still providing the reliability necessary.

The question of reliability is a question of how much damage is caused by a failing component or by someone gaining unauthorized access to a component. It is therefore closely related to security issues (cf. Section 9). However, no matter how reliable a component is there is always a chance of failure. It is very important that the remaining CAS doesn't fall into an undefined status once a component fails. The behavior of the remaining components must adapt to the failure and the network should stay operational. Hence, the question of reliability of the CAS is not only a question of reliability of its components but also a question of its adaptability.

If one assumes that such adaptability of the CAS is given the degree of reliability of its components may vary. Certain components will be more important to the ISP than others. For example a meter metering the traffic over an expansive transatlantic link will be much more important than a meter inside the ISP's network which gets used only for collecting statistic data (see Figure 7).




Figure 7: CAS Scenario - Reliability

This shows how the reliability of components is closely connected to their location. It is also closely connected to the replication of components because one way to make a system more reliable is to install several systems so that if one breaks down the others can do its job.

5.1.4 CAS Scenarios

After illustrating the useful distinction of the three dimensions of location, replication and reliability, three different combination options are discussed for the interconnection of two Access ISPs via a single Transit ISP.

Scenario (1) is depicted in Figure 8 and shows a situation, where two large Access ISPs are interconnected via a Transit ISP. The two Access ISPs serve a large number of end-customers, private ones or even corporate customers. For that reason at every router in

their network a high number of streams and lots of traffic will be entering the network. Therefore, the level of replication of components is based on these numbers and may require quite a lot of instances of the accounting component. Definitely, the border routers between each Access IPS and the Transit ISP are equipped with the full sequence of metering, accounting, and charge calculation components to handle a large number of flows and large traffic aggregates. The two Access ISPs must not show a complete similar situation, as depicted above. However, the main difference between the Access ISPs and the Transit ISP is defined by the fact that the Transit ISP will operate at higher speed, with larger traffic aggregates, and over longer distances. Therefore, the placement of components and their replication as well as their internal implementation depend on the actual requirements. Scenario (1) shows that two peering partner ISPs are connected to the Transit ISP domain at two different routers. At these ingress points the metering as well as the accounting component should be located for performance and operability reasons.

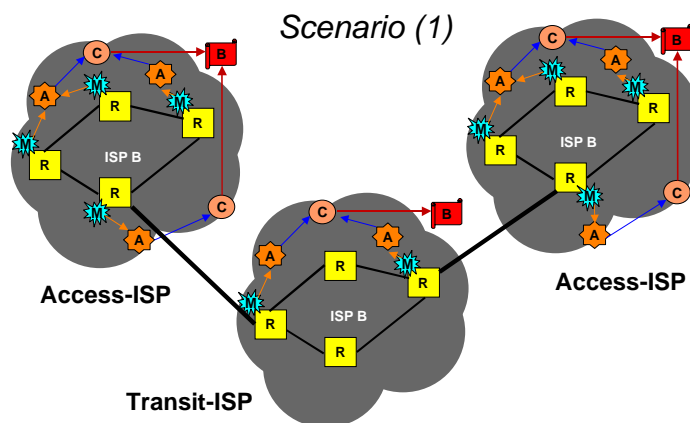


Figure 8: ISP Combination for Large Access ISPs (CAS Scenario1)

As depicted in Figure 8, the charge calculation component exists only once, which may be sufficient for some cases. Finally, the billing component is decentralized and dedicated to every single ISP in the market. It is out-sourced for the case presented, however, may be operated ISP-internal as well.

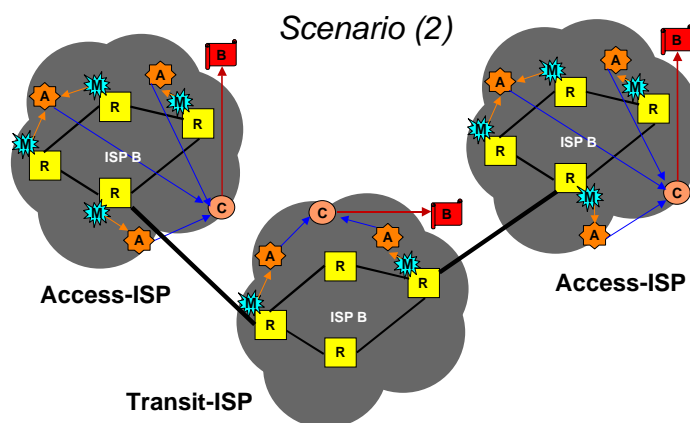


Figure 9: ISP Combination for Smaller Access ISPs (CAS Scenario 2)

Following along these arguments, Figure 9 outlines the scenario (2), where the number of end-customers is smaller compared to scenario (1), but the overall traffic transported per Access ISP is quite similar to scenario (1). This is reflected by the missing replication of the charge calculation component. Less end-customers to be served requires less

maintenance of charging records and financial data to be processed. Therefore, a single charge calculation component will be sufficient in terms of its performance. However, the technical components for handling accounting records and service-dependent QoS parameters are required in a similar scale. Billing services are kept in a similar fashion to scenario (1). Depending on the details, billing may be out-sourced from the Access ISP to the Transit ISP, once they operate in a close business relationship.

Finally, scenario (3) differs with respect to the billing service, which is completely out-sourced to a third party by one Access-ISP and the Transit-ISP. This may be a suitable solution for smaller Transit ISPs in a more local situation. All other technical features are similar to scenario (2).




Figure 10: ISP Combination for Smaller Access ISPs and Centralized Billing (CAS Scenario 3)

5.2 Discussion of CAS Dimensions

In the subsection above a set of possible scenarios has been presented. It identified several dimensions in which the components needed for implementing a CAS can vary. While these variations are explained for certain scenarios the discussion now considers general trends of these variations that can be observed. These trends can be important guidelines in designing a CAS for a specific ISP.

5.2.1 Location

The question of locating CAS components is a question of keeping them inhouse or outsourcing them. The one component that will always be kept inhouse is the metering. Metering is always done on the routers themselves or at least near to them (on links). It could only be out-sourced if the ISP out-sourced the network infrastructure itself.

Metering produces a lot of data. To transform this data into a form which can be used for further processing mediation is needed. One of the main aspects of mediation is that the output is much smaller than input. Therefore it is sensible to keep the mediation inhouse also because making the amount of data smaller is a precondition for outsourcing components. It is not sensible to transfer the original big amount of data to other entities for outsourcing.

Accounting is the first CAS component that realistically might be out-sourced. However the ISP might want to keep it inhouse to store the accounting records as an insurance against

fraud. Accounting is the first level where the amount of information is small enough to make it possible to store it over a longer period of time. It can later be used to check old bills.

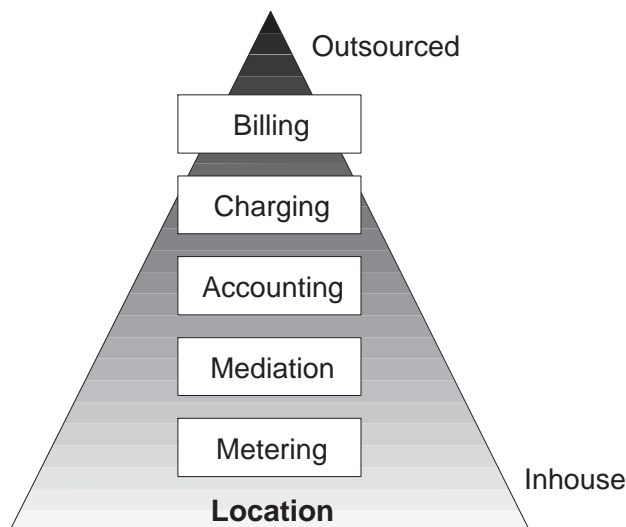


Figure 11: Sensibility of Outsourcing rises from Metering to Billing

Charging might well be out-sourced by smaller ISPs to companies which specialize in charging and pricing. These companies would provide the economic knowledge that small ISPs might not have.

The incentive to out-source billing is very high. As can be seen in Figure 11 billing involves communication with other providers. In the internet where many different providers can be involved in the provision of services this is not a trivial issue. Billing also involves dealing with customers and might even mean lawsuits with customers unwilling to pay. Therefore billing can be a big hindrance in the daily work of an ISP which is provision of services. It is very likely that only large ISPs will do their billing themselves. All other ISPs will out-source the billing to companies that specialize in this task (clearing houses).

So there is a tendency to keep lower level CAS components inhouse while outsourcing higher level ones. This is illustrated in Figure 11.

5.2.2 Replication

An important issue when designing a CAS for specific ISP is the number of units that is needed of each CAS component. Since an ISP normally has many routers and each router must be metered there will also be many metering units.

The metering units produce much data which must be preprocessed by mediation. Because of the amount of data it will be necessary in most cases to have as many mediation units as metering units. Only when the routers do not handle much network traffic and the mediation units are powerful can one mediation unit handle several metering units. However the data must still be transferred to the mediation unit. In general there will be as many mediation units as metering units and they will be very close to each other.

Accounting collects the data of customers. This means that data of several sessions will be stored in one place. Therefore and because the data coming out of mediation is much less there will be much less accounting units than metering and mediation units.

Charging gets the aggregated session data from accounting and decides how much the customer must pay for the session. There might be less charging units than accounting

units which would also be extremely useful for the necessary communication with the price setting component.

Billing will most probably be done by a central billing unit. It does not have to act on a short time scale and it is good to have all the billing in one place to keep control on it.

In general there is a tendency to have the more units the lower the layer of CAS is as is shown in Figure 12.

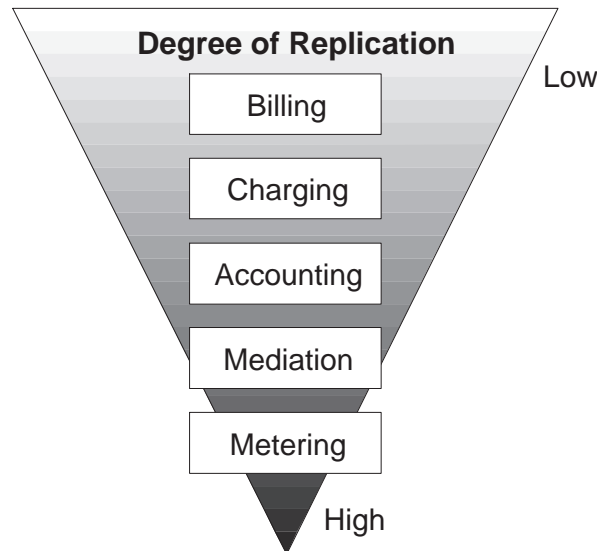


Figure 12: Degree of Replication decreases from Metering to Billing

5.2.3 Reliability

Each CAS component creates or processes information which is passed on to higher components or in case of billing to customers. Therefore, we argue that the importance of reliability of a CAS component corresponds to the importance of a single information record processed by this component.

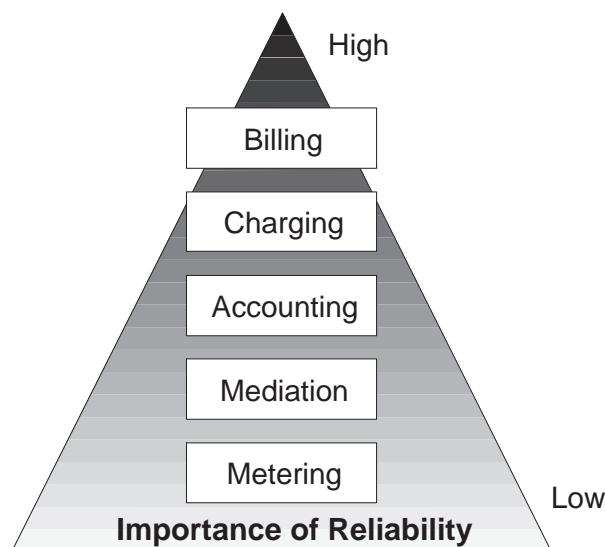


Figure 13: Importance of Reliability rises from Metering to Billing

Losing some metering data is probably the least important case. Losing a billing information record is the worst case, since a bill always includes data sent over a longer period of time,

e.g., a whole month. The same holds true for changing information records. It is only logical to assume that the importance of reliability of the CAS components rises from metering to billing as can be seen in Figure 13.

This has some major effects on CAS design. For one thing it means that security must be higher at higher CAS levels. The lower importance of information at metering level also implies that data has to be collected at bigger time scales and does not have to be 'exact to the byte'. However this isn't true anymore if dynamic prices due to congestions is intended. Laws might also make it necessary to do absolutely exact metering.

5.3 Service Interface and QoS Layer

Once understanding the task, characteristics, and dimensions of a CAS, its functionalities need to be designed within a distributed system fashion, allowing for, e.g., remote metering and local charge calculation. Therefore, these CAS functionalities as well as the service differentiation intended for multi-service IP networks need to be defined.

For the data path, the existing TCP/IP reference model lacks QoS expressiveness and hence does not suit the description of multi-service networks. In addition, for the control path in support of QoS as well as information required to be collected for the CAS, no scale has been defined until now.

Therefore, a new reference scale, which overcomes these shortcomings and which offers a platform for precise economic, architectural, and technical (implementation-oriented) notions has to be defined.

The traditional TCP/IP reference model is enhanced by inserting an explicit *Service Interface* between the Application and the Transport Layer. The Transport Layer is made "QoS expressive" by explicitly separating the data path from signaling concerns; subsequently it is called *QoS Layer* (cf. Figure 14). Furthermore, the reference scale implies a generic view of the IP Layer, which results in an accentuation of its technology independence, e.g., IPv4, IPv6, ATM layer, or IPX.

5.3.1 Service Interface Definition

Within Section 2 a service is specified as a particular task of an application, which just requires a network interaction to be performed. The Service Interface principally identifies and dissociates application requirements into single services. Furthermore, it provides access facilities to both either the QoS Layer or directly to the IP Layer. This necessity results from various service types¹ required by applications.

The Service Interface permits individual service configuration, maintenance, and supervision. It provides the necessary abstraction and transparency to applications and customers and makes no assumptions on technical service characteristics, e.g., if services are flows or aggregations.

With relation to M3I, and especially to the CAS, the Service Interface allows to design intuitive and comprehensible tools for customers. During a session a customer can instead of being charged per flow, stream, or packet volume, he can be charged per service. Consequently, an unprofessional customer is able to influence selectively and adapt

1. In this context *type* of service is used to distinguish services related to the QoS and the Application Layer, typically information services like DNS-lookups, WWW-browsing, and to services concerned with proper transport and synchronization, called infrastructure services, e.g., raw IP access.

precisely her application according to her budget. This is mainly a question of appropriate feedback.




Figure 14: Reference Scale and CAS Entity Placing

The customer influences the behavior of her application by configuring the Service Interface. This can be done automatically by the actually used application, e.g., by setting preferences, or by an explicit configuration of the Service Interface. The former method affects only the specific application whereas the later one affects the whole set of applications using the IP stack. A combination of both methods will be used in most cases, since the explicit configuration of the Service Layer corresponds to a general user policy and the application specific configuration corresponds to an adaption of it.

5.3.2 QoS Layer Definition

The comparison of where QoS entities, especially signaling, are placed nowadays in the TCP/IP reference model shows, that a large number of different and sometimes contradictory interpretations exist. Consequently the reference model has to be extended to fully enfold the concerns of QoS provision.

Roughly spoken the QoS Layer is a QoS enabled Transport Layer. It is divided in a *data part*, which is used for the provision of conventional data transport tasks, and a *control part*, which enfolds the subject of QoS and of signaling concerns. It has to be noted that all layers are involved within the provision of QoS and not only the QoS Layer¹.

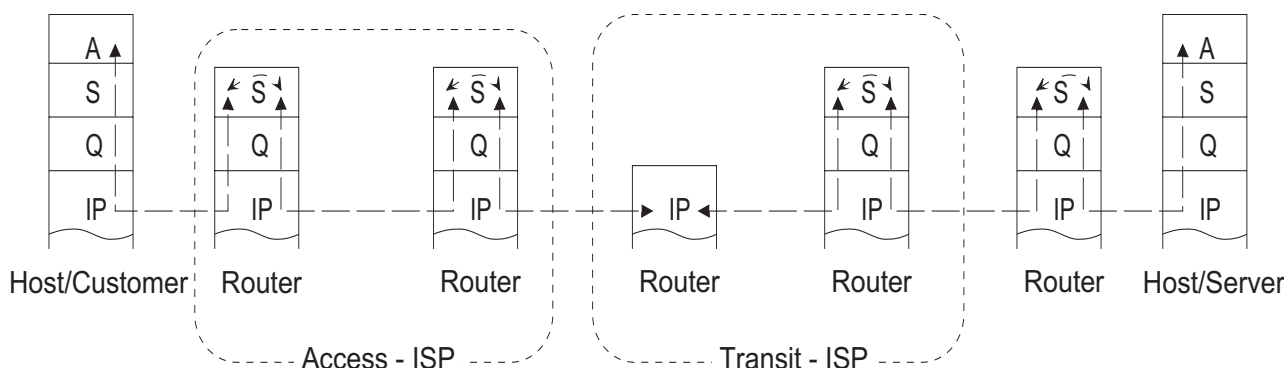


Figure 15: CAS Scenario: Customer - ISPs Communication

1. The Application Layer offers tools for intuitive QoS formulation, the Service Interface spreads and assigns the application's intuitive QoS onto particular services and the IP Layer applies and forces the adherence of the QoS parameters it has received from the QoS Layer.

The control part's primary task is to provide mapping functions from intuitive QoS suggestions into technically expressive parameters. The parameters take direct influence on the QoS Layer or are propagated and applied in the lower layers.

Concerning the CAS design, the QoS Layer represents the place, where metering, accounting, and charge calculation can be placed. These placements cannot be considered as strict, since the CAS entities affect parts of the IP Layer as well as parts of the Service Interface and the Application Layer (cf. Figure 14).

As shown in Figure 15, the reference scale has been transferred with slight modifications to the scenario as depicted in Figure 8. A single connection between a customer and an information end system is shown. The connection may represent a session or a part of a session, i.e. a single service. The focus is set on the layers required to express the functionality of the different routers. Depending on the router if it performs metering and/or mediation and/or accounting the number of layers varies. It has to be mentioned that this scenario does not cover the aspect of extra entities placed externally, e.g., on the wire, which manage the tasks of metering and accounting.

5.4 Mass Customization

This section outlines per customer dispersal in charging system design motivated by a perceived demand for mass customization in conjunction with distribution.¹ The context is intended to cover charging for a wealth of different services, but is particularly focused on network service given its inherent distribution. These ideas would over-complicate services that are not distributed, but are still relevant for non-distributed services, if they are often bundled with a distributed service, such as networking.

5.4.1 Motivation

Each customer is likely to have to be treated differently in some respect from other customers (Figure 16). There will be some cases, where multiple customers will be the same in all respects, but these will be exceptional co-incidences, rather than the rule. Differences will result from offering:

- A choice of tariffs.
- A choice of bundles with other products and services (whether other communication products such as mobile and fixed or higher level services such as TV packages).
- A choice of presentations of management and charging information (e.g., provider Web site, on board customer's own equipment, tailored for customer's accounting systems).
- A choice of aggregations of sites (e.g., all the mobiles, and all the personal use from within employer's networks as well as the fixed use from a family home; or all the branches of a business plus teleworkers).
- A choice of payment arrangements.

All this is beyond the usual choice of different levels of quality service, different availability, different usage patterns, which do not alter the structure of the systems needed to monitor

1. These ideas have been developed in BT since early 1996. They have generated considerable interest, but there has been no large-scale design based on them to date, only prototypes and demonstrators. Therefore, these ideas are untested in real scenarios that have evolved over time, rather than ones where the design can be created with the benefit of hindsight.

usage, only the results of monitoring it, and therefore are not relevant to mass customization of support systems.

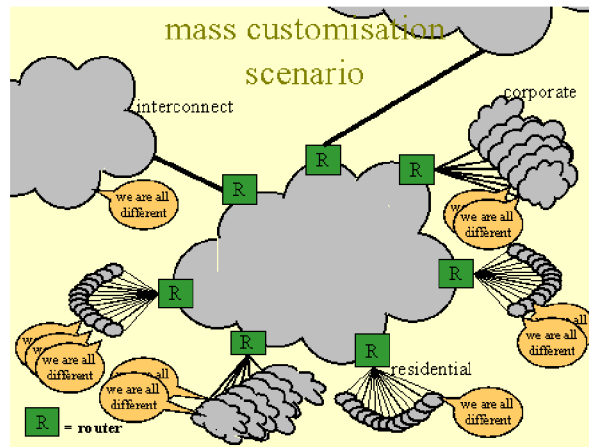


Figure 16: Customization Scenario

As well as mass customization, a related requirement is to allow new customers to bring their management data with them from their previous provider, and in a similar vein, allow them to take their data with them when they choose to move their custom to a competitor. Showing customers that they are not tied in is often perceived as a strong selling point, just as much as enabling them to import their relevant state from previous providers. However, structuring an architecture so that this is possible still allows a provider to deny the ability, or more subtly to partially deny it.

5.4.2 Structural Factors

Given each customer is more likely to be different from the next than not, each customer is assigned a logical instance of a Charging and Accounting System (CAS) (Figure 17). This instance is configured with policies specific to that customer. Initially M3I will not concern itself with the location or replication of these functions; it just assumes one per customer. Note that multi-site customers still have one CAS per customer, not one per site. Thus, the paramount objective is to be able to deal with customer activity in real-time, particularly if this involves collecting knowledge together.

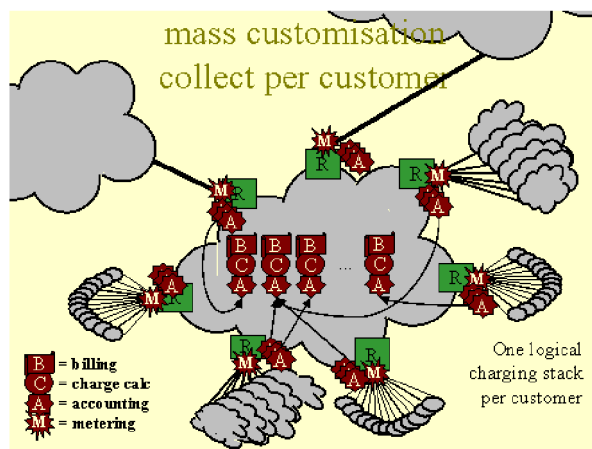


Figure 17: Collection per Customer

Clearly, there are features of each customer that are tied to topology, and cannot be divorced from location as the logical customer has been. For efficiency in dealing with large volumes of data, a meter must be placed on the service it measures. If this is the network service, it must be on the router, on customer premises, or somewhere on the link in between. If it is a video server, it must be monitoring the system log of the server.

Co-located with each meter are instances of 'stubs' of a mediation system, one per customer being measured at that meter. These stubs might do initial aggregation and also include the location of the rest of that customer's CAS. Thus, essentially they are rules in an object sitting locally at the output of the meter. The rest of each customer's mediation system sits with the logical CAS of that customer. Note that in Figure 17 the term accounting is used in place of mediation to align with the terminology of Section 2. However, mediation is used in this subsection, as it better expresses what is happening.

Having logically collected together the functions and data of each customer's CAS, one must separately choose where to locate the function and the data of each logical CAS (Figure 18).

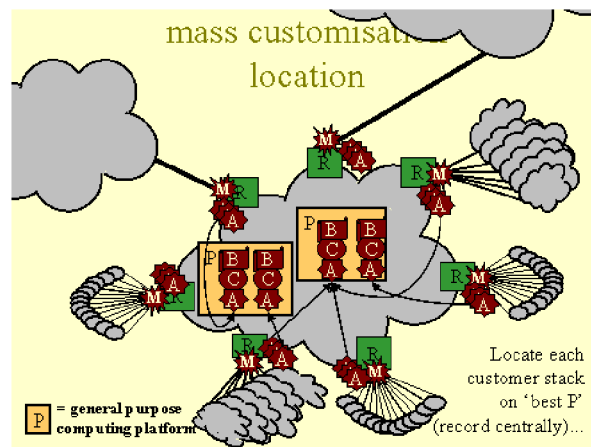


Figure 18: Location of Computing Platform

The decision may not only be where to locate the CAS within the provider's domain, but whether it makes sense to locate it in another party's domain. This requires taking into account many factors that limit charging & accounting system location.

Business factors:

- Strategic
- Cultural (e.g., "We've always done it this way!" or "whoever feels like operating it.")
- Contractual

Technical factors:

- Trust
- Efficiency (e.g., performance, timeliness, cost) of:
 - ◆ Accounting data transfer
 - ◆ Possible real-time interaction with customer accounting system
 - ◆ Policy distribution to the CAS (especially timeliness)
 - ◆ Relationship between data and process location (e.g., duplication of storage and/or processing steps)
 - ◆ Accessibility

- ◆ Reliability
- ◆ Including the advantages of statelessness (*e.g.*, reliability against failure)
- ◆ Cost of management of data/process evolution
- ◆ Cost of evolution from current legacy

Note that one of the reasons given for treating each customer's logical CAS separately is to allow extensions to the common schema required for certain specialist customers to be easier to develop without requiring initial central design authority. That is, 'stove-pipes' can be re-assimilated into the common design as they become common, rather than trying to suppress them. This is the sort of factor requiring taking into account under the heading of evolution costs.

Also note that there are no specialized computing resources required for a CAS - it merely needs processing, storage and communication - in other words a partition of the resources of any general computing platform. All common charging processes are assumed to no longer require specialized archaic machinery like bill printers or envelope stuffers. However, those that do, have to include this as a factor in their location. Only at this point should a customer's CAS be distributed across multiple platforms, and then only, if there is no single best platform location, *e.g.*, for multi-site customers where the data feeds need to be real-time and therefore are difficult to aggregate.

Figure 20 illustrates the one of the bullet points above - interaction with the customer's accounting system. As already stated, this will be one of the factors determining the best location of the provider's logical CAS for each customer. Not all customers will have their own accounting system, but those that do will most-likely require their Internet-related accounts to be included in their book-keeping, which can be done in real-time, just as many of the other more aspects of each customer's own business will be reckoned in real-time. Note that a customer view of its own financial status is also necessary for the customer's price reaction function [21].




Figure 19: Bespoke Interfaces

Finally, Figure 19 shows the provider-centric function of its accounting system. Nothing at this stage needs to be accurate in real-time. It is sufficient to collect the information from each customer's CAS for analysis on a provider-wide basis on more relaxed time-scales, *i.e.* in batch, perhaps daily.

Collection across all the provider's systems is necessary for maintaining the company revenue status, marketing analysis and other, general data warehousing operations. Although each customer-centric CAS can be tailored to each customer, each must also

implement a few common functions necessary to interface with the central requirements of the provider's data warehousing functions. Further, it should be possible to locate each customer's CAS using a common mechanism.




Figure 20: Customer First, Provider Last

The intention is to free the service offered to each customer from the 'lowest common denominator' constraints of a monolithic design, without creating a panoply of 'smoke-stack' designs that become a nightmare to maintain. In practice most of the functions of a CAS are usually implemented over a relational database. Therefore, the skill is to design an extensible schema where any commonality between customizations can be exploited, rather than leading to divergent schema for each customer type.

5.4.3 Mass Customization Conclusions

The traditional centralized telecommunication accounting system has been turned inside-out, being instead primarily structured to serve the information and management needs of each customer and only secondarily those of the provider.

5.5 CAS Inter-provider Architecture

While Section 3 presented the general network of WP 5 and 6 it is important for the design of the CAS to take a closer look at the components and interfaces of this architecture which are directly connected to the CAS. Of special importance are the connections between components of the same type which were not shown in Figure 3.

Figure 21 shows the interaction which takes place between two ISPs when two of their customers exchange data. The figure does not show that several other ISPs might be positioned between these two providers, but this can easily be imagined and does not change the general situation. The components and interfaces presented in the figure are those with direct relevance for the design of the CAS. Components will be presented in more detail in Section 6 and interfaces in Section 7.

Interaction between the two providers takes place on two levels. The first one is of course on the data path since providers must exchange data between their networks. Since the transport of this data is not free ISPs will charge each other for the data transported. This leads directly to the second level of interaction between them. Each provider collects information about the transported amount of data and calculates a fee for it. He later sends the responsible entity a bill through a billing system. He either bills another provider if he

provided a service for him or sends the bill to one of his customers. In this case he will get a bill from another ISP if he used the other ISP's service to give a service to his customer.

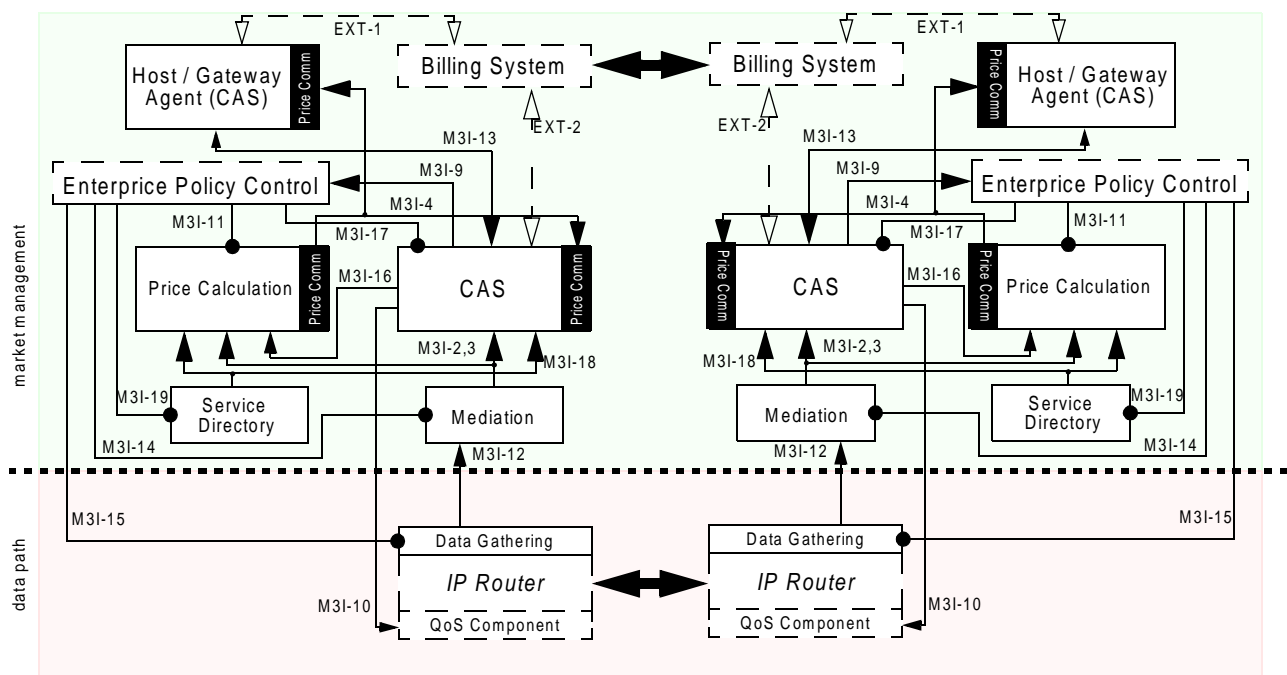


Figure 21: Conceptual view of a Charging and Accounting System

Information exchange between providers occurs on the level of the billing systems, where inter-provider invoices are exchanged. Instead of performing absolute billing between interconnected providers, they can also offset their claims against each other. A huge set of peering agreements and settlement schemes exist for today Internet Service Providers, however, they are defined in a quite static manner and do not allow for immediate responses to bandwidth bottlenecks or further customer and user demands. Further inter-provider information exchange happens as part of specific protocol processing as defined in the QoS model applied, e.g., for resource reservation purposes such as using the Resource Reservation Protocol RSVP or inter-Bandwidth-Broker communication, where messages are exchanged between the boundary routers of neighboring providers. In any case, a type of signaling or consolidation protocol to take care of the distributed information scattered around in the network of networks is required.

6 CAS Components

Based on the overall architecture for M3I and the embedded CAS within the components and interfaces presented in Figure 3 and Figure 21, the basic building blocks surrounding the CAS are shown and the main tasks of the CAS internals are presented in Section 6.3.

6.1 Data Gathering (Metering)

In Figure 20 Data Gathering is integrated in the IP router. Alternatively it could be placed directly on the wire. Indeed such a solution introduces supplementary expenditures (entity needs its own IP address¹, requires probably special protocols), furthermore it can only

1. The IP address is needed since the metering/data gathering entity needs to be configured according to the Enterprise Policy.

supervise the actual usage of the link and has no knowledge of the usage of the critical and for congestion control relevant resources within the router. The interconnection of several gathering units to reconstruct the current router status is not a feasible approach. Concluding from this, it would in spite of having metering units on the wire, be necessary to know the state of the router, so an explicit interaction of the CAS and the routers would anyway be required.

Functions

In [21] the data gathering entity's tasks has been divided into the three main parts of measuring, sampling and local mediation. There is no need to refine them. The only extension to be made concerns the categorization of the traffic and the interface to the mediation entity.

Among other things in Section 4.1 the *support for transport, service and content charging* has been declared as the objective of the CAS system. Basically, Metering can be seen as the foundation of the CAS. Nevertheless, it is only able to support transport metering and partially the metering of services. Content charging has to be integrated by other means.

In this context, transport charging and thus metering of transport characteristics is mainly interpreted as the measurement of raw data volume, whereas the metering of service characteristics additionally comprehends QoS relevant measurement, i.e., particularly focusing on time critical assurances.

The interface to the mediation entity will be explained later.

Implementation Choices

- Netflow Router.
- Linux machines with kernel modifications.

6.2 Mediation

A definition of mediation can be found in Section 2. The purpose of the mediation entity is thus as described there, to transform metered data (of each single meter), to merge the data of different meters and thus to reduce the amount of data. The order of the transformation and the merging is not yet determined, but can be relevant when taking into account performance aspects.

The effect that mediation units dramatically reduce the amount of data has a deep impact on network design. Strategies where to place mediation stringently need to be evolved.

Functions

- Receive metering data (from one metering unit).
- Transform metering data (e.g. into collected information referring to one customer).
- Send transformed data to Accounting.

Implementation Choices

- Own code (very flexible).
- SIU [12] [13] (advantage of scalability to large networks, less implementation effort).

6.3 Charging and Accounting System

The Charging and Accounting System entity consists of a charging, an accounting, a customer support and a user support component. The separation of the CAS into components increases the degree of freedom, since the components can be physically distributed. The embedding of the CAS into the overall M3I architecture is provided through eight interfaces.

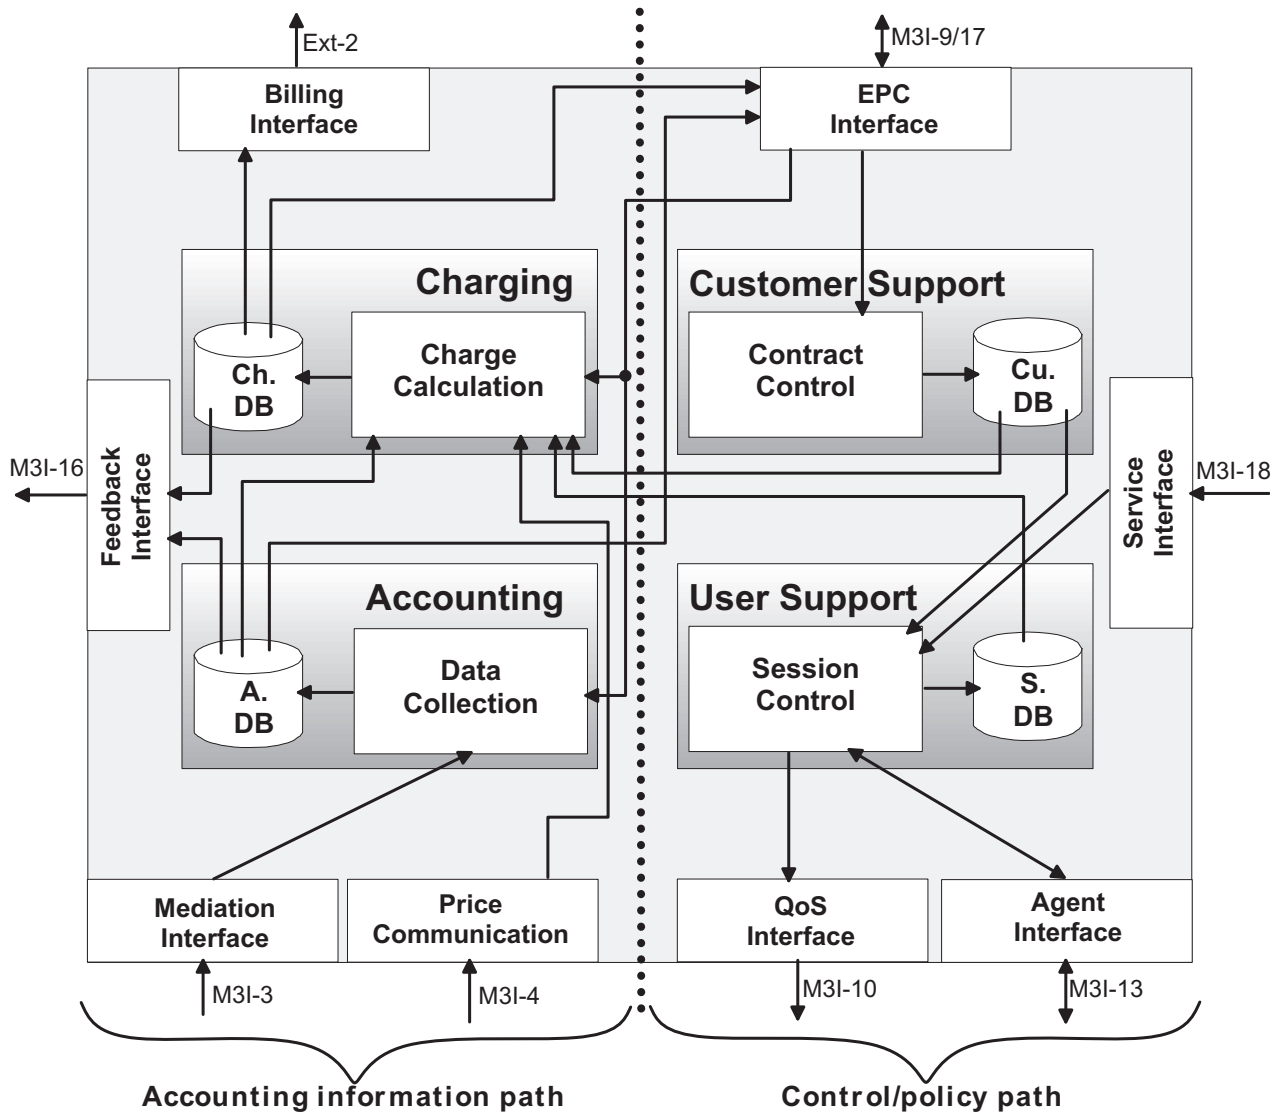


Figure 22: The Charging and Accounting System in Detail

The CAS can be divided in two logical paths as shown in Figure 21. One, the Accounting Information Path, depicts the flow of the charging data. The other, called Control/Policy Path, is used to manage and configure the CAS, especially the entities involved with the processing of the charging data. The two paths differ in the direction they process. The Accounting Information Path starts from the bottom and ends on the top, whereas the Control/Policy Path starts from the top and processes down to the bottom.

6.3.1 Accounting

The Accounting component gets the metered and mediated usage data and is responsible for storing it. It must provide this stored data to other components and interfaces for further

processing, feedback or statistic evaluation. Accounting is the central usage data storage component.

Functions

- Receiving the transformed metering data from Mediation.
- Merging the mediation data from several mediation units into complete accounting records for each customer. This is done in Data Collection.
- Storing accounting records for future use and as insurance against fraud. The Accounting Database is used for the storing.
- Send accounting records to Charging. This isn't done in an active way. Instead the Accounting Database is accessed by the Charge Calculation to retrieve the required data.

Implementation Choices

- SIU + Database.
- Own code + Database.

6.3.2 Charging

Charging processes the mediated and metered usage data further. It calculates the appropriate charges for the resource usage using a tariff. To be able to do this it needs input from various other components like Price Calculation and User Support.

Functions

- Receive accounting records from Accounting by accessing the Accounting Database. It is still an open issue when this is done or how the data is pushed out of the database to Charge Calculation.
- Receive prices for different services from Price Setting. Price Communication is used for this.
- Receive information about services from Service Support by accessing the service database.
- Calculate charges using the received prices, the retrieved accounting records and the service information about the used service.
- Calculate possible discounts.
- Collect charging records for a billing period. The charging records are stored in the Charging Database.
- Send charging records to Billing. This is not done in an active way. Instead billing can access the Charging Database and retrieve the required charging records. This is done periodically, e.g., once a month for each customer.

Implementation Choices

- Own code + Database.

6.3.3 Customer Support

As stated in Section 5.4 an ISP can have many and a large number of different customers. Additionally, a customer is not the same as a user, e.g., one customer might pay the bills of several users. The customer is the entity which has a contract with the ISP. The contents of

this contract, e.g., number of users covered by the contract and their names and accounts, are managed within the Customer Support.

Functions

- Store contract with customers.
- Forward relevant contract information (e.g., users covered by the contract, services allowed) to User Support on demand.
- Forward relevant contract information (e.g., discounts) to Charging on demand.

Implementation Choices

- Own code.

6.3.4 User Support

While the Customer Support is responsible for keeping all contract information the User Support is responsible for making sure that those contracts are kept. On the one hand this means that he blocks any user requests that are not covered by the contract the user (more exactly: the customer which pays for the user which may not be the user himself) has. On the other hand he must make sure that a service requested by a user is delivered to him if the contract allows it.

Functions

- Receive requests for services from host agents and negotiate a service with them. This is done in session control. To clarify whether a user is entitled to a certain service session control may access the customer database. The session information is stored in the session database.
- Make sure that agreed on services are provided. This is done by reading the QoS requirements for the service (service definition) out of the service directory and passing them to the QoS component of the router.
- Pass service definitions to Charging on demand so that Charging can calculate charges for a utilized service. This is done by forwarding the appropriate session data from the session database.

Implementation Choices

- Own code.

6.3.5 Interfaces

These are the interfaces to the other M3I components that the CAS must exchange data with. In some cases the interfaces might be part of CAS components. However, they are all shown as separate components in Figure 22 because at the current state it is not yet clear which interfaces will be included and which will stay separate. In general it will be useful to separate the interfaces when they are communicating with several other components. This can be several components of the same type (e.g., the Mediation interface will receive data from several mediation components) or of different type (e.g., the EPM interface communicates with several components inside the CAS). It is also sensible to create a separate interface whenever this interface might be reused in other parts of M3I (e.g., the price communication).

-
- QoS interface: To provide services to the customer it is necessary to control the QoS component of routers. This interface can be used to set the QoS parameters of routers.
 - Mediation interface: The Mediation interface is responsible for collecting data from several Mediation interfaces, possibly even from Mediation interfaces of different types.
 - Enterprise Policy Control interface: This is the interface for changing parameters of the CAS after the system has been deployed. By using this interface the EPC can install new services or request and receive charging or accounting data.
 - Service Interface: This interface can be used to read service definitions out of the Service Directory.
 - Billing Interface: This interface is responsible for sending the calculated charging records to the Billing System.
 - Price Communication: Price Calculation is responsible for setting the prices used by Charging. To send the calculated prices to Charging the Price Communication is used.
 - Feedback Interface: To set prices the Price Calculation uses price models with various input variables. Some price models need usage or charge information as input variables, hence these can be communicated to the Price Calculation via the Feedback Interface.
 - Agent interface: This interface is responsible for communication with the customer. Right now this mainly includes the selection of services the customer can use. However the functionality of this interface might be enhanced at a later stage, since it is likely that there is need for more communication between the customer and the CAS.

6.4 Price Calculation

Price Calculation entity is extensively documented in [18].

6.5 Billing System

Although the Billing System exceeds the scope of duties of M3I, some assumptions on its functioning need to be made. The knowledge of the Billing System affects the CAS in so far that the interfaces can be placed correctly and that plausible and complete scenarios can be created. The most important assumption on the Billing System is mentioned in Section 5.3, where it is used to describe inter-provider charging concerns.

Functions

- Collects charging records on a per user basis and for a longer period of time (e.g., a month). This can be done by selecting the needed information from the Charging Database.
- Can access a user database which contains user ID, address, personal information, etc.
- Creates bills for each customer.
- Sends bills to customers and receives payment.

Implementation Choices

- Not in M3I
- High Probability of outsourcing to clearing house.

6.6 Enterprise Policy Control

The Enterprise Policy Control entity represents the network provider's interface for the management and supervision of all (except the Billing System) CAS related entities. The design of the Enterprise Policy Control entity will not be illustrated in the M3I project.

Functions

- Provides an administration interface to various components.
- Receives data from the CAS as a response of the actions taken.
- Can change price setting parameters and models.
- Can change data to be metered and how it is mediated to CAS and Price Setting.
- *E.g.*, provides usage statistics, congestion warnings, error notifications.

Implementation Choices

- Not in M3I, some interfaces to the CAS components have to be created, though.

6.7 Host/Gateway Agent

The Host/Gateway Agent has two different functions. The first one is to communicate charges to hosts and gateways. Apart from charge communication a host agent acts on behalf of the user. This can include negotiating services with the CAS or automatic reaction to the communicated charges even payment. A host agent can also restrict a user's options when the customer wants to control the behavior of the users he pays for.

Functions

- Act on behalf of the user.
- Include customer functions.
- Negotiate a service with the CAS.
- Act on feedback from CAS, *e.g.*, charge communication during a session.

Implementation Choices

- Own code.

6.8 Service Directory

As has been mentioned before a user has no real understanding of quality of service in technical terms. He is unable to specify his requirements in a way that can be used as a direct input to QoS technologies. Instead he has a more higher level view of quality. This higher level view must be translated into technical values which can be used for the setting of QoS components in the network and for charging according to the technical usage data. This translation takes place in the Service Directory.

Functions

- Keep a list of available services and their QoS requirements. This list is can be accessed by Price Setting so that prices for services can be set. The list can also be accessed by the CAS which is responsible for setting the QoS components on routers according to the service definitions.

- Allow changing of service->QoS requirement mapping and installation of new services. This is done by allowing the Enterprise Policy Control to access the service directory.

Implementation Choices

- Own code.

7 CAS Interfaces

The interfaces required for the communication between the various components of the CAS and the environment as defined in the architecture (cf. Figure 3) are enlisted in Table 1. In general, these interfaces are designed to act as protocols, allowing for the communication between two remote entities of these components, or as software interfaces, reflecting the clear architectural decision, that the interaction between those components happens within a common address space.

Number	Type	Entities	M3I WPs
M3I-3	protocol	Mediation - CAS	5.1, 6
M3I-4	protocol	Price Communication - Price Communication	5.2, (6)
M3I-5	software	CAS - Price Communication	5.2, 6
M3I-6	software	Price Setting - Price Communication	5.1, 5.2
M3I-9	protocol	CAS - Enterprise Policy Control	6
M3I-10	protocol	CAS - QoS Component	6
M3I-11	protocol	Enterprise Policy Control - Price Setting	4.2, 4.3, 5.1, 6
M3I-12	protocol	Data Gathering - Mediation	5.1, 6
M3I-13	protocol	CAS - CAS	6
M3I-14	protocol	Enterprise Policy Control - Mediation	5.1, 6
M3I-15	protocol	Enterprise Policy Control - Data Gathering	5.1, 6
M3I-16	protocol	CAS - Price Setting	5.2, 6

Table 1: System Interfaces

7.1 General Interface Design Criteria

The definition of interfaces for component interactions need to obey a number of criteria to allow for the consideration of the overall CAS requirements as presented in Section 4. In particular, the interfacing efficiency is of utmost importance. Firstly, this implies that the data structure developed for these interfaces and protocols need to be context-free (*i.e.*, well-defined portions of the protocol data units need to be located at identical places). This will ease the performance-efficient implementation of the protocol at a later stage.

Secondly, the processing requirements for the protocol itself are based on the data structure defined as well as on the finite state machine (FSM) designed. These FSMs should be as minimal as possible and need to minimize the number of state transitions for processing incoming protocol data units.

Thirdly, the memory consumption of protocol data structures as well as the protocol entity should be minimized as well. However, this will be mainly an essential design criterion for protocol entities only, which will be placed on routers only. Any server or host-based system may not delimit this criteria.

Finally, the bandwidth utilized for exchanging these data structures across networks should be minimal as well. However, due to the processing optimizations in the first place, the pure protocol-dependent data structure will be a second dimension of optimization only.

For all protocols and interfaces the time scale of interactions and feedback is essential. Depending on per packet-feedback (such as ECN marks, TCP flow-control, or forwarding decisions taken inside routers) in a very small time-scale of milliseconds, control loops and decision trees need to be designed. Larger time scales will exist on a price and tariff communication layer, where minutes (e.g., the duration of a session) or some hours will occur for application-to-application control. For ISP-to-ISP control future Service Level Agreements (SLA) show a quite similar time-scale in the lower range, ranging from hours to months. Hence, service provisioning and network planning determine the larger time-scales, not being considered particularly within M3I. But the data aggregated and collected by the CAS may be used as input signals. Finally, business strategies in the largest time-scale form the upper end of interactions, ranging from days to months.




Figure 23: Relevant Time-scales for the Design of Interfaces for the CAS

7.2 Mediation - CAS (M3I-3)

The usage data, which has been mediated after the data gathering took place, needs to be transferred to the CAS. Therefore, a protocol has to be developed, which defines the rules and transmission units for transferring mediated data to the CAS component. Since the anticipated load for this interface probably will be high, the protocol must be highly efficient, yet extensible.

The data exchanged across this interface will include one of the following alternatives, which depend on the particular scenario:

- A simple hand-over of data gathered by Data Gathering.
- A hand-over of data mediated based on the particular inputs from Enterprise Policy Control. This may result in the dedicated specification of specialized data to be required for the CAS, some special aggregation of these data, or even the neglecting of data resulting from the gathering process.

One particular aspect of M3I is given by the investigation, whether M3I-2 (cf. [19] for details) and M3I-3 can be designed as a single protocol. Due to the fact that the type of data for M3I-2 and M3I-3 are similar, only differing in its size, the number of data records, or its frequency of exchange, a similar type of protocol is envisioned. Further details will be

presented during the fine-grained design of the pricing tools design as well as of the CAS itself.

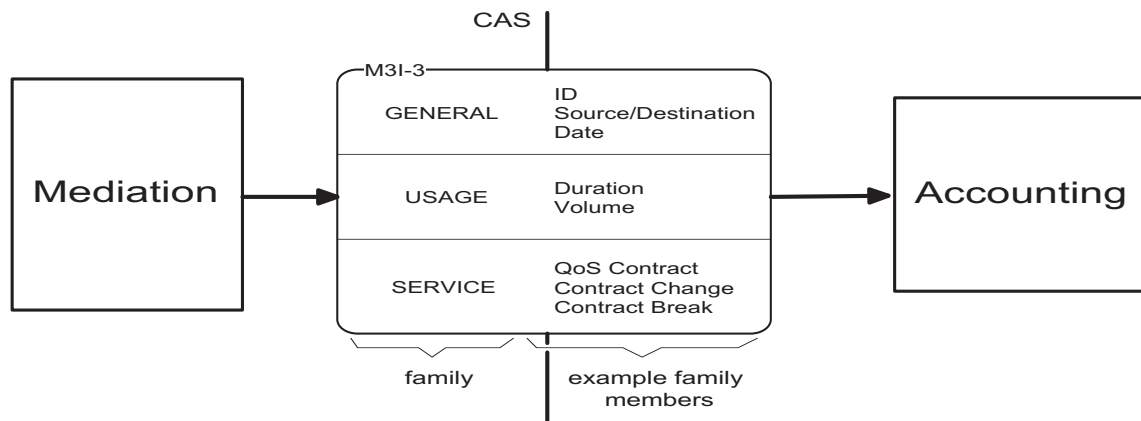


Figure 24: Mediation Interface, M3I-3

Family	Member	Description
General	ID	Identification of customer and session (or service).
	Source/Destination	Originator and Consumer identification.
	Date	Date of service/session allocation.
Usage	Duration	Duration of service/session use ^a .
	Volume	Effective bandwidth consumption.
Service	QoS Contract	Agreements on jitter, loss rate, delay etc..
	Contract Change	Customer initiated renegotiation.
	Contract Break	Notification of network failures ^b .

Table 2: Mediation Interface families

a. The duration will probably correspond to the rate in which the mediation unit sends its collected records to the accounting entity. In this case it does not comply with the duration of the entire service use.
 b. E.g., almost all resources in the network are sold and suddenly a link goes down. There will be no possibility to continue to keep the contracts.

7.3 CAS - Price Communication (M3I-4)

The CAS requires input from the price setting component, in particular details on prices and tariffs for services. A generic way to exchange this price information between different participants has to be defined, and M3I-4 specifies the protocol to exchange prices and tariffs. The final protocol (or protocol elements, depending on the fine design process) has to allow for the widest scope of information, yet efficient transmission. Further details are presented in Section 5.5 of [19].

CAS-relevant details for this protocol encompass the following ones:

- Memory-efficient data structure for prices and tariffs, based on basic thoughts expressed in [2].
- Processing-efficient data structures (context-free) for prices and tariffs.

- Inclusion of customer and/or user identification information for prices and tariffs.

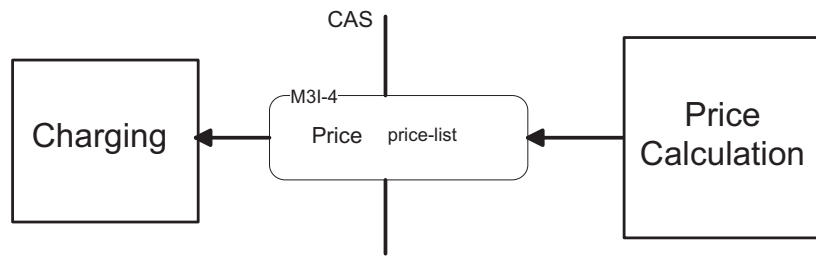


Figure 25: CAS-Price Communication Interface M3I-4

Since the price communication does not reflect a full and self-existing component as such, rather a well-defined means for communicating prices and tariffs across networks and between components, the software-based interfaces between the “hosting” component and the price communication are of local matter for the components only.

Interface	Price	Description
M3I-4	Price-list	List of price/unit, e.g., per volume, per jitter and delay constraints etc. Prices may vary depending on time of day, congestion, etc.

Table 3: M3I-4 Interface

7.4 CAS - Enterprise Policy Control (M3I-9/M3I-17)

This interface for exchanging information for Enterprise Policy Control touches the edge of the M3I project work. However, for establishing a full-fledged cycle of feedbacks between the parties involved in running an ISP business, this interface determines the feedback loop to the business strategies. Based on the information, processed data gathered at routers and mediated in-between, a usage or service profile of a given (sub-)network can be obtained. Therefore, these data present important decision basics for overall business strategies.

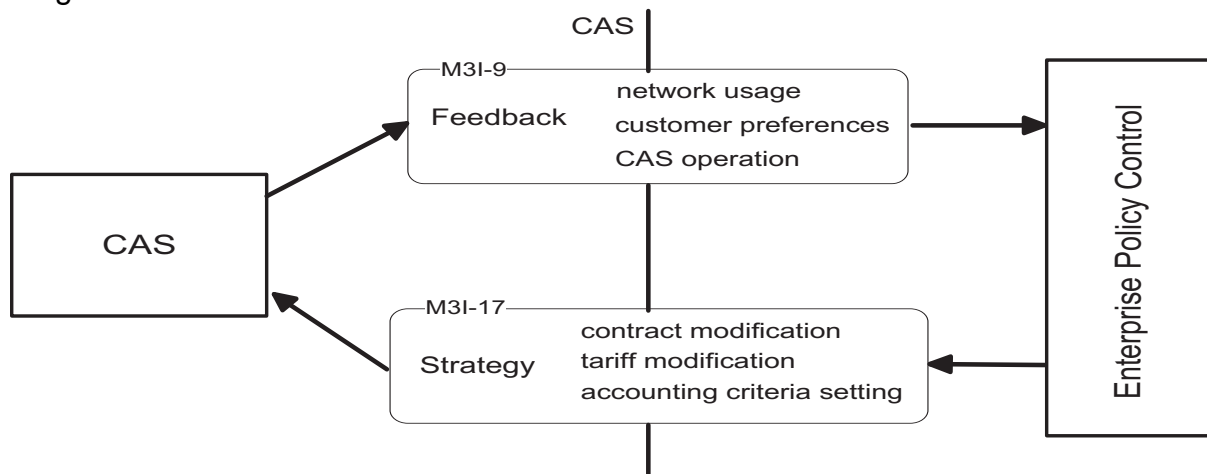


Figure 26: CAS-Enterprise Policy Control Interfaces M3I-9 and M3I-17

The relevant protocol data structures and the type of interaction required between the CAS and Enterprise Policy Control are beyond the scope of WP6's work. However, an open and extensible protocol is envisioned to be provided for the prototypical system's design.

Interface	Feedback/Strategy	Description
M3I-9	Network usage	Quantitative and spatio-temporal usage of resources, statistical vs current usage situation.
	Customer preferences	Evaluation of user preferences for creation of user profiles and contracts.
	CAS operation	Supervision of the CAS.
M3I-17	Service profile changing	Creation of new, removal and adaption of existing contracts as a reaction to new market strategies.
	Tariff modification	Adaption of the tariff algorithms to the new or modified contracts and user profiles.
	Accounting criteria setting	Configuration of the accounting entity according to new or modified contracts.

Table 4: M3I-9 and M3I-17 Interface Functionality

7.5 CAS - QoS Component (M3I-10)

This protocol interface allows for the reconfiguration of the QoS component resulting from tasks being performed within the CAS. This could be the reconfiguration of the ECN marking algorithm or admission and access control on various timescales as described in [1]. It can also be used to make sure that a specific service agreed on with a customer is provided. While the customer's interest is that his application behaves the way he wants it to he isn't interested in detailed matters of data forwarding through the network. Therefore, he agrees on a specific service with his ISP and the ISP must translate this service to detailed QoS requirements. To fulfil these requirements QoS parameters in each router can be set by using the M3I-10 protocol. This can be done using Flowspecs or other measures.

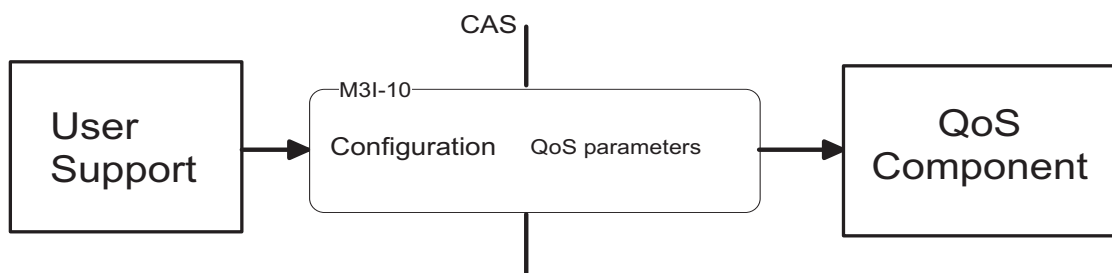


Figure 27: CAS-QoS Component Interface M3I-10

Interface	Configuration	Description
M3I-10	QoS parameters	The QoS parameters needed to provide a service. Taken from the definition of the service in the Service Directory

Table 5: M3I-10 Interface

7.6 CAS - Host/Gateway Agent (M3I-13)

The CAS-CAS protocol defines the charge communication required for delivering data to host/gateway agents of the CAS. As depicted in Figure 3, to this end, charge

communication is currently expressed as a regular interface. However, there are certain similarities expected between price communication (cf. [19]) and charge communication. Thereby, charge communication can be considered and designed as distributed middleware, as well, providing transparent transport of charging information over a variety of transmission channels. As with price communication, a large fraction of work is likely to be spent for designing appropriate protocol data structures as well as specifying a reasonable interface for the access of charge information. There is a potential for optimization through alignment of this interface with price communication, though, in case of service signalling.




Figure 28: CAS-Host/Gateway Agent Interface M3I-13

Interface	Negotiation	Description
M3I-13	Service request	A service request by a user.
	Service response	The acknowledgement of a request, e.g., acceptance, denial, alternative service suggestion.
	Charges	The charge feedback to the user or to a gateway of the customer.

Table 6: M3I-13 Interface

However, the CAS-Host/Gateway Agent interface is used for more purposes than only for charge communication. Every communication between users or customers and the CAS of their ISP is handled by this interface. This especially includes the negotiation of services. Users send service requests to the CAS if they want to utilize a certain service. The CAS uses the interface to transmit its response to the user, which can either be an acceptance or a denial of the service request.

7.7 Feedback Interface (M3I-16)

The feedback interface provides the necessary input for the Price Calculation entity to adapt and change the prices according to the current network status. The feedback represents the status of the network in terms of resource usage, i.e. quantitative usage

given by the amount and characteristic of the accounting and charging records, as well as in terms of the current charging and accounting policy.




Figure 29: Feedback Interface M3I-16

Interface	Feedback	Description
M3I-16	Accounting templates	Information about the characteristic of the Accounting Records.
	Charging templates	Information about the tariffs and the resulting Charging Records.
	Charging and accounting volumes	Quantitative feedback about accounting and charging ^a .

Table 7: M3I-16 Interface

a. The qualitative feedback is given by the accounting and charging templates

7.8 Service Interface (M3I-18)

The designed CAS is strongly service orientated. Users can request services and get charged for the use of these services. To provide a service QoS components of routers are set according to service definitions. Therefore, the CAS must be able to access these definitions which are stored in the Service Directory. The purpose of the service interface M3I-18 is to send the needed service definitions from the Service Directory to the CAS.

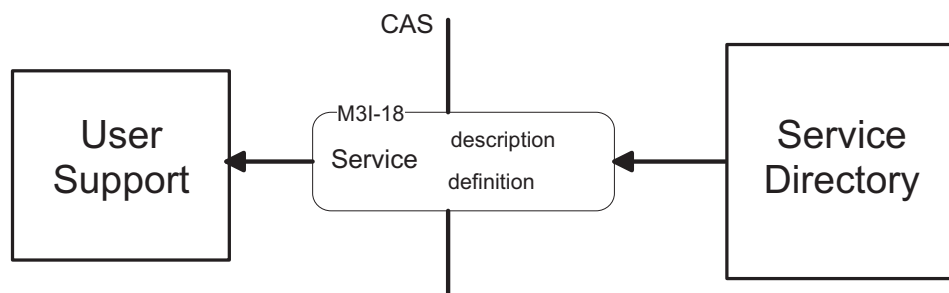


Figure 30: Service Interface M3I-18

Interface	Service	Description
M3I-18	Description	The general description of a service as seen by the user.
	Definition	The technical definition of a service in terms of QoS parameters.

Table 8: M3I-18 Interface

7.9 CAS - Billing Center Interface (Ext-2)

Ext-2 is not only an interface between CAS and Billing Center, it is also the interface that the M3I consortium provides to a possible billing institutions.

Ext-2 supports facilities to clearly identify the customer and the enterprise, which is responsible for the accounting. Furthermore, it identifies the offered service. This is done by giving it an ID¹ and by attaching the characteristics of the service. The service contract and the costs are as well passed to the interface. The contract can be used by the Billing Center to legitimate and detail list the costs.




Figure 31: CAS - Billing Center Interface Ext-2

Billing Information	Information Detail	Description
ID	Customer	ID of customer.
	Enterprise	ID of the enterprise providing the policy control.
Service	Identification	Identification and characteristics (volume, duration, date etc.) of the service.
	Contract	Agreements at service initiation, changes of contract, effective provided service (statistical, average values).
	Costs	The costs the customer has to pay. summation vs detailed listing of costs.

Table 9: Ext-2 Interface

8 Service Example

Applying the defined components and interfaces to a given application scenario allows for an initial description of an example and its practicability. Therefore, Figure 32 sketches the case of a video supervision application. Imagine a hospital or university supervising its geographically distributed buildings in a centralized way.

1.The ID corresponds to one in the database of the Charging entity.

The video supervision application provides a simple interface to the user. The data propagated downstream, i.e. to the video cameras, is splitted into single services, e.g., the handling of each single video channel could be splitted into several 'independent' services. Upstream data is simply merged and passed to the application.

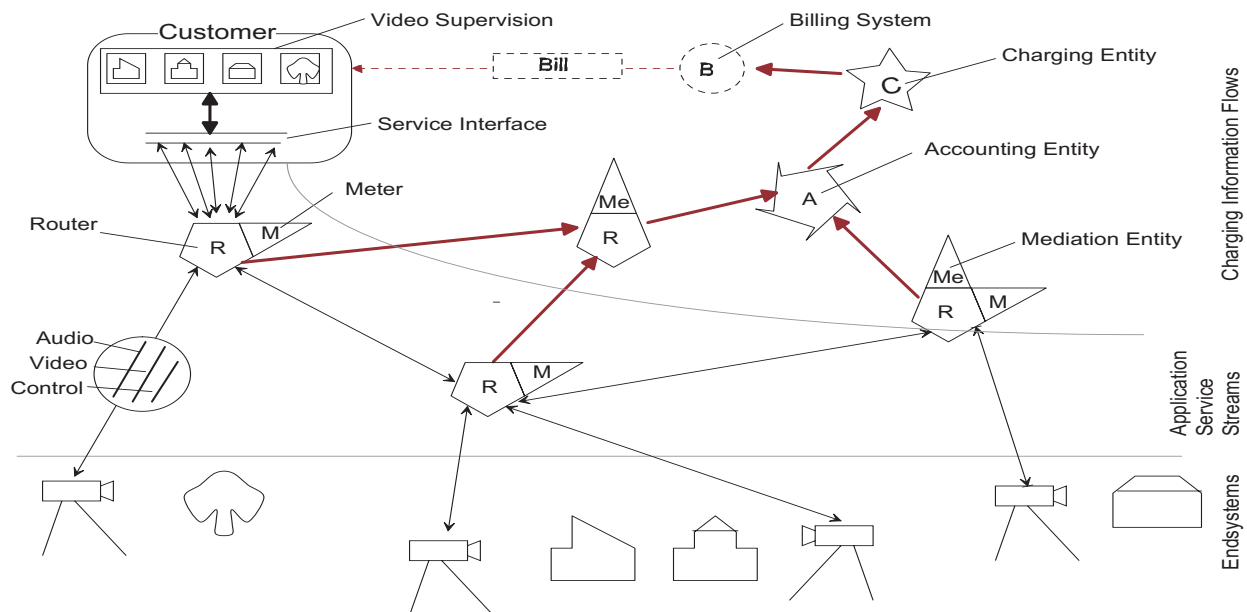


Figure 32: Video Supervision

The user can influence the quality and the costs of his application by changing QoS parameters. This can be done by setting application preferences or/and by configuring the Service Interface. However a configuration of the Service Interface automatically affects all applications accessing the network. Thus the configuration of the Service Layer sets a general preference and the configuration of application preferences adapts them according to the needs of the current application.

The splitting of the application requirements into services can be the prerequisite for user friendly cost control. In the example it is presumed that the user/customer can see the costs that each single camera generates, since the CAS charges and accounts per service. Note that each video channel consists of three services, an audio, a video and a control service. The knowledge of the costs each service generates is an important feedback and permits to precisely change QoS parameters (e.g., color - b/w, resolution, voice: on/off), resulting in an optimal cost control to quality ratio.

What can further be seen in the example is the tentatively placing of the CAS entities. Most routers only meter the video services. The metered data probably consists of a session id, a user id, a time of day, and a date, as well as of the offered and the negotiated Type-of-Service (ToS).

The metered records of the single video services are periodically transferred to a mediation capable router. This one selects and aggregates the metered information (e.g., the user ids, the time the service is used etc. need not to be associated anymore with each record). After the aggregation, the mediated data is sent to Accounting and later on to the charging unit. With the information received from the Charging entity the Billing System is able to generate a bill, which is then handed over to the customer.

8.1 Example Interaction between CAS and Price Calculation

Given a dedicated networking scenario, including the available signaling and policy distribution protocols, on one hand the particular combination of the interfaces M3I-12, M3I-3, and M3I-10 as well as their interaction modes can resemble a Common Open Policy Service (COPS) scenario. On the other hand, it can be optimized for low-latency propagation of administrative admission decisions.

Concerning the exchange of information between the two components of the CAS and the Price Calculation all objects of charging-related information, other than authentication, should be considered optional in order to limit computational complexity for the default case. Second, it is important that charging-related handling of these service invocations can be delegated to another entity within CAS. Therefore, the architecture is in-line with that one developed in the IETF RAP working group [36]. In fact, the aim is to fit the whole pricing communication design in the framework developed there. One of the main advantages of such a decoupled architecture is that the filtering process of incoming service requests is transparent for the outside, because service requests can be redirected to (or intercepted by) the edge router, a separate service enabler, or yet another policy server. Each of these components can delegate decisions or instruct another component with the results of a local decision. As long as the service interface is not affected by such internal decisions, any combination can be employed, which creates the highest flexibility, for example, to choose whether pricing and charging information is tied with service information or transmitted separately.

Some examples for possible M3I considerations may encounter the following ones:

- Auction models, while bids are piggybacked onto RSVP.
- Usage-based pricing, counting all or certain packets or ECN marks.

9 CAS Security Considerations

Security issues arising during the design of the CAS are dealt by in a more centralized fashion. For a productive CAS, secure transactions and all information exchanges have to accompany CAS processes required to perform the designed tasks. Various confidential information is exchanged between the parties of all accounting, charging, and billing components, which need to be protected from malicious attackers. According to [14], [15], and [16], the development of a security architecture consists of several steps.

The first step performs the risk analysis by identifying assets, threats, and vulnerabilities before evaluating the risk of an attack by combining the probability of an occurring incident and its impact. The second step designs the concept for a secure solution before the concept will be implemented in a third step. Therefore, the basic security requirements to be met are described in the following. Additionally, important assets and threats are identified before security measures are introduced to prevent malicious attacks.

9.1 Considered Security Requirements

In general, a set of five widely accepted security requirements have to be met by the design of the security architecture for the CAS.

Confidentiality – A confidential information exchange determined that only the sending and the receiving CAS component are able to interpret the data exchanged. Any other

entity not involved in the process between these two CAS components only sees the encrypted data, where no interpretation is possible.

Integrity – The integrity of data exchanged is guaranteed, if the data cannot be changed unnoticed during the transfer.

Non-repudiation – Neither the sending nor the receiving CAS component can deny that a specific message has been sent/received. Both the sending and the receiving component can prove that a specific message with the corresponding content has been sent and received, respectively. This requirement founds the basis for any type of auditing, which may become important to clear inter-provider or provider/customer complications.

Authenticity – The authenticity is the proof of the identity of a communicating component. In general, authenticity can be achieved by proofing possession (e.g., passport/ID), knowledge (e.g., password/PIN), or existence (e.g., biometrical password) of a unique feature of identification.

Availability – Finally, a system is available, if it works permanently under any pre-defined conditions. Every state of the system needs to keep all necessary conditions of a pre-defined functional and performance profile.

In addition, the question of safety is not part of this work, as the overall M3I system will not address them either.

Anticipating the detailed CAS design as well as the overall M3I architecture, the achievements of these security requirements need to be detailed in advance. In particular, as soon as strict performance guidelines and functionality decompositions in addition to high level security requirements are articulated, a trade-off between these issues have to be taken. The description of assets and threats will allow for a qualitative judgement, which trade-off can be optimized for performance reasons and which trade-off will focus on higher degrees of security.

Furthermore, these security considerations assume that system-level security for end-systems, hosts, routers, and required devices is given. Therefore, the following assets and threats are focussing on the exchange of data and information required to be performed between the defined components within Section 6.

9.2 CAS-oriented Assets and Threats

According to the M3I requirements discussed and presented in [1], every single business entity represents a number of assets that could be attacked. In addition, any type of local information kept within such a business entity could be intercepted by a malicious third party. Based on the basic roles identified, the following list comprises important assets and threats being present in the set of M3I scenarios. Some roles, even though termed differently, are equivalent to the ones defined in [25]. Other roles as presented in [1] are added for completing the views on M3I roles for Internet stakeholders. These assets can be reflected in the valuation of service quality sold to customers. Therefore, the flow of technical data representing this valuation, finally expressed in money or monetary equivalents, is data, which needs to be metered, collected, accounted for, and charged for by the CAS.

9.2.1 Consumer (Customer)

Consumers use services from providers. A single end-consumer (a person) or a corporation may act as customers. Therefore, end-customer and enterprises need to be distinguished for security reasons.

9.2.2 Provider

The stakeholder termed provider offers services to consumers. However, since the type of service offered, the service utilization, or the maintenance of this service may vary, providers need to be separated out by further means which influence the degree of security required.

In addition, the physical networking infrastructure utilized within the Internet determines the technological environment, in which the threats may occur.

9.2.3 End-customer

End-customers want to consume services, in particular, they may want to perform a type of electronic business with a provider offering a product or service. *E.g.*, in order to reserve bandwidth from an end-customer to a service provider for a dedicated network connection, the end-customer has to establish a link by means of his end-user network provider, possibly the access provider, and possibly a backbone provider. Once this type of network connection has been agreed upon, it is required for the end-customer to be able to access this network connection always, to receive the Quality-of-Service (QoS) reserved for this connection. On one hand, he must be made aware of the fact that he needs to pay for the claimed services, if this has been negotiated or advertised in advance. On the other hand, he can assume safely, that the negotiated services are provided within the contract limits agreed upon.

9.2.4 End-user Network Provider (Customer Premises Network)

According to the Internet market layer model [1], it is possible that end-users¹ are affiliated to an End-user Network Provider, commonly known as Customer Premises Network (CPN). This network may offer additional local applications or dedicated conditions for data transport as well as general services. In this case, the CPN has to make sure that these services are always available for end-customers (reflecting the business scope) and that data is transported according to the security requirements to all end-customers. Since CPNs tend to be operated and maintained by a single authority within a given enterprise organization, this domain has to follow enterprise-wide security strategies. Denial-of-Service attacks may originate from end-users (not necessarily end-customers, as hackers do not have a business relationship with the CPN) within the CPN or from users crossing the inter-connected Access Provider. In case of remote access operations by mobile end-users, appropriate mechanisms have to be provided.

9.2.5 Access Provider

The Access Provider is responsible for setting up the network connection between the end-customer and any type of Information Provider or the Backbone Provider according to pre-

1. The terminology applied in the requirements document of M3I [1] does not distinguish between user and customer. However, for the CAS design this distinction is relevant, since the term customer reflects the business relationship with a provider (a.o., being concerned with the flow of money) and the user reflects the technical relationship with a provider, utilizing a service or running an application.

defined parameters. It needs to make sure that it is always available for end-customers to perform reservations, to utilize services, or to quote prices as long as the initial contract has negotiated these issues. The Access Provider, as the Backbone Provider or the End-user Network Provider as well, must protect himself from denial-of-service attacks and he also needs to establish or maintain his business relationships with other providers (e.g., Backbone Providers or any type of Information Provider based on his business models) by dynamically adapting the negotiated Service Level Agreements (SLA). An Access Provider also needs to make sure that he receives the payment from the end-customer for the set of provided services.

9.2.6 Backbone Provider

A Backbone Provider runs a backbone network, which inter-connects other backbones or Access Providers. Its main task is to provide availability of the backbone services (mainly connectivity) and transport data packets reliably through its network according to the previously defined QoS parameters and SLAs. As for Access Providers, it needs to make sure that it receives payments according to the services provided.

9.2.7 Data Center Provider

Depending on the particular business model, the Data Center Provider has to support various server technologies for Information Providers. While on one hand the interconnectivity to a Backbone Provider is essential (its security assets and threats are similar to the ones for Access and Backbone Providers), on the other hand the reliability and availability of servers is business critical.

9.2.8 Market Place Provider

Security considerations for the information providers are quite similar and will not be part of the basic M3I work. Therefore in brief, the Market Place Provider needs to ensure that his platform is an open one, accessible by anybody who is allowed to do so, a secure one, and a reliable one. Otherwise, business opportunities will migrate to business losses. There exists related work on the provisioning of these type of market places elsewhere.

9.2.9 Communication Service Provider

The Communication Service Provider seems to be a bit different in this sense (being part of information providers), since it offers services, which require the direct use of an underlying communication infrastructure. Therefore, the network layer security provided by an Access Provider, the End-user Network Provider, and the Backbone Provider determine the level of security achievable. In case this degree is not sufficient, an application level security can be provided and integrated by the Communication Service Provider, since they offer services such as fax, e-mail, or packet-based audio and these services are implemented by dedicated services (traditionally termed applications).

9.2.10 Content Provider

As mentioned above for the market place provider, content as such can be secured in various fashions and various tools. Details should be obtained in related work elsewhere.

9.2.11 Application Service Provider

For the Application Service Provider as similar statement holds as for the Communication Service Provider, even though the type of application considered here is different. Certainly,

application level security mechanisms are required to bridge the gaps of network level security offered by the Connectivity Providers.

9.3 Security Measures

To counter threats and protect assets, different security measures have to be taken into account. An important distinction needs to be done first. The user data flow has to be separated from control data (management and maintenance data), since different security requirements apply to both of them. While user data may be transmitted with a certain tolerable loss rate (*e.g.*, for videoconferencing or surveillance services), control data always has to be transmitted confidentially, since it contains control information for the communication itself, the network setup process, or private charging and accounting information for end-customers. Unlike for these control data, the level of confidentiality for the user data flow needs to be defined individually by the end-customer through QoS parameters. The CAS itself may be involved in this distinction, since secured control data acceptance and transfer may be considered a value added service provided. In this case, appropriate means to account for this service are required.

A second distinction includes the subdivision of communications into different kinds of associations and flows depending on the connected business entities. *E.g.*, the connection between an end-customer and his Access Provider meets different security requirements than the connection between two competing Backbone Providers. The CAS needs to identify the types of flows exchanged between customers to allow for the detection of QoS and its automated charging.

9.3.1 Protection of Confidentiality

To guarantee a high level of confidentiality, the information exchanged should be encrypted with a secret session key that is exchanged regularly to avoid replay attacks. However, the open question remains for the level of granularity for which data have to be encrypted during the exchange between CAS components. *E.g.*, data metered at a single router port will be mediated and accounted for based on the enterprise policy. How important are these data and for how long do they need to be secured? Certainly, the accounting records generated are significantly more in danger than the pure metering data, since already a particular level of aggregation has taken place. Charging records constructed out of these accounting data need to be secured even further, since usage information has been correlated with customer identification, and a malicious interception would allow for the generation of customer profiles, *e.g.*, in terms of service usage, information access, or partners.

Possibilities for a secret key exchange are the Diffie-Hellman algorithm or various asymmetric cryptographic algorithms [28]. Within the Internet environment, the IPsec protocol suite has been proposed [24]. It comprises a set of standards to provide confidentiality at the IP layer, which forms an important charging layer as well. Its IP Encapsulating Security Payload (ESP) provides encryption for IPv4 and IPv6 packets and would be helpful for CAS inter-component communication, such as the collection of accounting records or the distribution of bills. However, the degree of confidentiality achieved will be based on the decisions taken for the implementation of the protocols between CAS and PM components.

9.3.2 Protection of Integrity

Data integrity can be implemented using hash-functions. A hash-function calculates the hash-value (or message digest) of the original message. This value represents a fingerprint of the message and it allows for the guarantee of the integrity of the sent data, since it is consistent, collision-free, and one-way [28]. To make sure that this hash-value and the message was not replaced by a malicious attacker, the hash-value needs to be transmitted in addition to the message, either on a secured channel or included in a digital signature.

A digital signature could consist out of the hash-value encrypted by the sender's asymmetric private key. The IPsec Authentication Header (AH) commonly uses a keyed hash-function rather than a digital signature, since digital signature technology is too slow and greatly reduces network throughput [28].

In particular for the CAS, the set of components interacting together for allowing for the preparing of the final invoice, is limited and its identities may be storable, including the keying material in case of a public scheme. However, it needs reconfiguration as soon as new CAS components will be integrated into a given network infrastructure.

9.3.3 Protection of Availability

Hardware solutions to protect the availability of a network include redundancy in hardware components. To avoid a loss of data or content when a system component fails, static information should always be stored on more than one disk (e.g., in a so called RAID System - Redundant Array of Inexpensive Disks). One way to save dynamic information on a system is to mirror all actions performed on a running unit on a "backup unit". In case the running system crashes, the backup-system can take over without or with minimal time delay. Software solutions include the monitoring of performed actions in a log file enabling the recovery of the last consistent state (e.g., commit or rollback in database systems).

For the CAS design, the redundancy question of components and their interactions plays the central role. As discussed in previous sections, the replication degree of CAS components depends heavily on the physical topology of the network and their interconnection links. In addition, the availability degree to be achieved determines a second influence on the location and replication questions discussed above.

9.3.4 Protection of Non-repudiation

The concept of non-repudiation of origin (NRO) proofs to a receiver that a message has been sent by a sender. Digital signatures are one way to guarantee the non-repudiation of origin. In addition, non-repudiation of delivery (NRD) proofs to a sender, that a message has been received by a receiver. One way to guarantee non-repudiation of delivery is to return a digitally signed acknowledgment of the receipt to the sender [25].

In the optimal case, the CAS has to deal with the support of legal aspects regulated from the government. Therefore, at least at a higher level of aggregation, e.g., for an invoice, NRD is essential for the bill delivery. However, as M3I can not tackle legal aspects of data communications with great depth, these non-repudiation protections will not be considered in greater detail.

9.3.5 Protection of Authenticity

The safest way to authenticate with passwords is to use one-time passwords that change whenever two parties want to authenticate. E.g., the S/KEY authentication program is based on hash-functions and avoids replay-attacks [28]. A digital signature consists of the

hash-value encrypted by the sender's asymmetric private key. The receiver can authenticate the sender, if he is able to decrypt the message digest with the sender's asymmetric public key, since only the sender should have the according asymmetric private key.

For the CAS the customer's authenticity is the main concern, in particular, for the case of assigning charging records to customers. On one hand, to ensure that customers are correctly identified and, on the other hand, that data arriving from a CAS are authenticated. Furthermore, metered data originating from routers needs to be taken care of, to ensure that it really originates from the locations assumed. Otherwise, input data to Mediation as well as other components may be faked.

9.3.6 Auditing

The process of auditing requires many of the above presented concepts and mechanisms being in place. In particular, the data storage of information being accounted for forms the basic problem, in traditional telecommunications acts need to be applied in the Internet communications domain. However, as the M3I work will not focus on legal acts and regulatory telecommunications laws, these problems will not be discussed.

Concerning the technical auditing possibilities, e.g., between two ISPs, a commonly agreed upon set of parameters and values, most suitably based on Service Level Agreements negotiated in-advance, need to be maintained for all ongoing communications. Due to the large amount of data considered, the auditing will be limited to certain average value calculations, some dedicated peak communication times, or some statistical measurement methods of traffic at specialized network interfaces. Certainly, the reliability as well as ensured integrity of these data are the prerequisites for a successful auditing.

10 Related Work

The area of related work with respect to the CAS covers a set of different view points from which the significant ones are presented. For the Internet-oriented systems the terminology applied and architecture options developed are presented. In addition, the view point of CASHMAN is included as well to provide a comparison from the ATM perspective.

10.1 Terminology

The definitions of components and tasks as well as their naming form the basis for further comparisons and discussions. To allow for a good overview on terminology applied in selected approaches of the charging work, critical terms are listed.

10.1.1 Charging

Charging is doubtlessly, besides accounting, the most important terms in the domain of the CAS. Based on the Webster's Dictionary [35] "to charge" is explained, a.o., as "to impose or record a financial obligation".

Standards and research work tend to show a quite close understanding of the tasks and definition for charging. Comparing the following citations lead to a quite common basis, which is reflected in the charging definition applied for the M3I CAS (cf. Section 2).

The European Telecommunications Standardization Institute ETSI [7] defines as follows: “Charging is the determination of the charge units to be assigned to the service utilization (i.e. the usage of chargeable related elements).”

Within [22] a full process point of view is defined as follows: “Once these accounting records are collected and prices are determined in full pricing schemes on unit service, e.g., encompassing different quality levels for services or service bundles, the data for an invoice need to be calculated. The process of this calculation is termed charge calculation, performing the application of prices of unit services onto accounted for records determining the resource consumption. Thus, the charging function transforms mathematically unequivocal technical parameter values into monetary units. These units need to be collected, if they appear at different locations in the given networking environment, and are stored in charging records. Of course, accounting as well as charging records determine a critical set of data which need to be secured to ensure its integrity when applied to calculate monetary values or when used to compute an invoice’s total.”

[33] defines: “Charging determines the process of calculating the cost of a resource by using a price for a given accounting record, which determines a particular resource consumption. Thus, charging defines a function which translates technical values into monetary units. The monetary charging information is included in charging records. Prices already may be available for particular resources in the accounting record or any suitable resource combination depending on the network technology or application.”

The charging process for business models offering ATM services is also termed “rating and discounting process” [30] and is “responsible for the charge calculation according to a specific pricing policy and using the collected usage data.” Therefore, charging mechanisms correlate service usage and calculate the charge the customer is faced with after the service utilization.

Finally, [5] describes: “Charging is the process of evaluating costs for usage of resources. Different cost metrics may be applied to the same usage of resources, and may be allocated in parallel. An example would be a detailed evaluation of resource consumption for further processing by the service provider, and a simple evaluation of resource usage for online display of current costs. A detailed evaluation of the resource consumption can be used for generating bills to the customer, or for internal analysis by the service provider. A simple evaluation of current costs can be used for displaying an estimation of accumulated costs for the service user, or for control purposes by the customer organization or by the provider. Cost allocation assigns costs to specific endpoints, such as sender and receivers of a multicast group.”

10.1.2 Accounting

Accounting is the other extremely important terms in the domain of the CAS. Based on the Webster’s Dictionary [35] “accounting” is explained as “the system of recording and summarizing business and financial transactions and analyzing, verifying, and reporting the results”.

While the charging term tends to be used quite similarly, accounting shows two different points of view. The first one is related to economic theory, where accounting relates to business processes, including profits and benefits. The second one relates to technical aspects, where technical parameters are measured and collected.

ETSI [7] defines accounting as the “revenue sharing amongst operators”. The ITU [18] distinguishes further terms: “Accounting revenue division procedure: The procedure

whereby accounting revenue is shared between terminal administrations and, as appropriate, between the administrations of transit countries". "Accounting rate: The rate agreed between administrations in a given relation that is used for the establishment of international accounts". "Accounting rate share: The part of the accounting rate corresponding to the facilities made available in each country; this share is fixed by agreement among the Administrations."

A technical explanation on the tasks and interfaces for accounting is presented in [22]: "Therefore, these units need to be accounted for, traditionally performed on a per-call basis over time. However, in packet-switched networks, the accounted for information may encounter a huge number of different parameters, e.g., number of packets sent, duration of a communication, number of transactions performed, distance of the communication peer, number of hops traversed, or bandwidth used. Depending on the protocol layer applied for this accounting task, only a subset of accounted for parameters are useful. In general the accounting record determines the container for collecting this information. These records and their special appearances depend on the networking technology used, such as N-ISDN, ATM, Frame Relay, or IP. They can also be created for application services, for example, the call data record is being used for this purposes in H.323 IP telephony. Further, the Real-time Flow Measurement working group within the IETF investigates appropriate accounting mechanisms."

[33] defines as follows: "Accounting determines the collection of information in relation to a customer's service utilization being expressed in resource usage or consumption. Thus, accounting defines a functions from a particular resource usage into technical values. The information to be collected is determined by a parameter set included within an accounting record. This record depends on (1) the network infrastructure, which supports the service, e.g., Internet, N-ISDN, ATM, or Frame Relay, and (2) the service provided. The content of an accounting record is of technical nature, such as the duration of a phone call, the distance of a high-speed network link utilized, or the number of market transactions done. This accounting record forms the basis for charging and billing."

The accounting process applied to ATM services is defined in [30] and complies with the ITU-T process definitions summarized above from [18].

Finally, [5] outlines: "The process of accounting involves the following functions: collection of usage data by usage meters, creation of accounting records (data structures, or protocol data units of an accounting protocol), transport of accounting records, and collection of usage data by an accounting server."

10.1.3 Further Terms

Interfacing the CAS requires to accept, react, and offer certain messages and information. Therefore, the terms **metering**, **pricing**, and **billing** are summarized at this stage as well.

A broad commonality and conformance can be observed for metering. ETSI [7] defines Metering as "[...] the measurement of 'components' which can be used for charging such as the duration of the call [...] named also 'collection of charging information'." A full task and term definition for metering is included in [22]: "[...] there remains a single technical prerequisite for identifying and collecting accounting data. This process is called metering. Based on existing technical equipment in operation, the metering tasks identify the technical value of a given resource and determine their current usage. If possible, metering can be tied to signalling events. Otherwise, it may be performed regularly, e.g., every ten seconds or every hour, it may be stimulated on other external events, such as polling requests, or it may be performed according to some statistical sampling scheme. In that

case, it is closely related to network monitoring. The IETF's Management Information Bases (MIB) for switched networks and the Simple Network Management Protocol (SNMP) architectural framework may provide a means of keeping monitored data." Also for the ATM approach, network element usage metering functions are described, being responsible for the generation and reporting of accountable resource information [30].

A quite similar use of pricing has been observed with respect to related work. ETSI [7] defines pricing as "[...] the correlation between 'money' and 'goods' or 'service'," while it is noted that "the term is not generally used in telecommunications, the usual term being 'tariffing'." [33] says: "Pricing is the process of setting a price on a service, a product, or on content. This process is an integral and critical part of businesses and closely related to marketing."

Finally, billing denotes the "[...] process of transferring the stored charging information for a user into a bill" [7]. This is in close relation to "the process of consolidating charging records on a per customer basis and delivering a certain aggregate of these records to a customer is termed billing" [22] as well as with [5] and [33]. [30] distinguishes between various billing mechanisms and options based on the form of the bill (e.g., itemized or aggregated) or the time of delivery (e.g., periodic, per-call, or pre-paid).

10.2 Systems

The number of projects concerned with charging and accounting tasks in the Internet increased quite significantly over time. Therefore, only a number of recent and charging-centric work of system's design and modeling is summarized below. Another and a more detailed overview can be found in [31].

10.2.1 The CATI Project

The objectives of the Swiss National Science Foundation project CATI (Charging and Accounting Technology for the Internet) [32] included the design, implementation, and evaluation of charging and accounting mechanisms for Internet services and Virtual Private Networks. This covered the enabling technology support for open, Internet-based Electronic Commerce platforms in terms of usage-based transport service charging as well as high-quality Internet transport services and its advanced and flexible configurations for VPNs. In addition, security-relevant and trust-related issues in charging, accounting, and billing processes have been investigated. Important application scenarios, such as an Internet telephony application, demonstrated the applicability and efficiency of the developed approaches. This work was complemented by investigations of cost recovery for Internet Service Providers, including various investigations of suitable usage-sensitive pricing models for end-to-end communications based on reservations as well as Service Level Agreements in-between service providers.

10.2.2 Charging Internet Services

Many projects dealing with charging and accounting functionality on the network level try to achieve a high independence from pricing models [33]. However, it has been articulated that pricing in general and usage-based pricing in particular can impose a high overhead on telecommunication systems [20], [29]. Any form of usage-based pricing for various telecommunication services is interesting, because underlying resources (such as satellites, frequencies, cables, routers/switches, and most notable operating personnel) are scarce and very costly. The traditional Internet pricing model has been critiqued constantly

in the past years for its economic draw-backs of not being incentive-compatible [29], [6], and [10]. Furthermore, it is inflexible — for example, it does not allow for combined sender/receiver payments — and does not provide economic signals which are needed for network planning and expansion. But most importantly, the current model is based on the assumption of a single service best-effort network that provides a similar service to all customers. Therefore, the multi-service paradigm needs to be investigated with respect to heterogeneous networking infrastructures and technologies of the Internet.

10.2.3 Lightweight Policing and Charging

The main assumption of this work is that a multi-service packet network may be achieved by adding classification and scheduling to routers, but not policing [4]. Therefore, a lightweight, packet-granularity charging system has been investigated emulating a highly open policing function, which is separated from the data path. The amount of charging functions required depends on the customer's selection of services and is operated on the customer's platform. The proposed architecture includes a set of functions distributed to customers, which may include metering, accounting, and billing as well as per-packet or per-flow policing and admission control. The proposal concludes that lower cost is achieved through simplicity without sacrificing commercial flexibility or security. Different reasons for charging, such as inter-provider charging, multicast charging, and open bundling of network charges with those for higher class services, are all catered for within the same design.

10.2.4 A Role Model for Charging

For Intelligent Networks (IN), charging issues have been considered to allow for a clear view on a separate charging service on its own right or as a part of the overall architecture [27]. The developed role model is part of the service framework to support the creation of telecommunication services by utilizing re-usable components. A service constituent defines reusable components for building services, including parameters and mechanisms, and is independent of any particular service. The charging constituent has been defined, which is applicable to public telecommunication services, but is independent of them. The role model itself offers five roles (client, charge handler, customer profile handler, charge calculator, and database handler), each of which being concerned with charging-related tasks. For these roles the message exchanges required to exchange charging information as well as the stimuli occurring in this model have been investigated.

11 References

- [1] R. Andreassen (Edt.): *Requirements Specifications, Part I Reference Model*; M3I Deliverable 1; Version 5, June 13, 2000.
- [2] B. Briscoe (Edt.): *Architecture, Part I Primitives & Compositions*; M3I Deliverable 2; Version 1, June 12, 2000
- [3] B. Briscoe: *The Direction of Value Flow in Connectionless Networks*; First International Workshop on Networked Group Communication (NGC'99), Pisa, Italy, November 17-20, 1999
- [4] B. Briscoe, M. Rizzo, J. Tassel, K. Damianakis: *Lightweight Policing and Charging for Packet Networks*; Third IEEE Conference on Open Architectures and Network Programming (OpenArch 2000), Tel Aviv, Israel, March 26-27, 2000

-
- [5] G. Carle, M. Smirnov, T. Zseby: *Charging and Accounting Architectures for IP Multicast Integrated Services over ATM*; 4th International Symposium on Interworking (Interworking'98), Ottawa, Canada, July 1998.
 - [6] D. Clark: *Combining Sender and Receiver Payments in the Internet*, Telecommunications Research Policy Conference, October 1996.
 - [7] ETSI: *Internet Protocol (IP) based Networks; Parameters and Mechanisms for CHarging*; ETSI TR 101 734 V.1.1.1, Sophia Antipolis, France, September 1999.
 - [8] G. Fankhauser, B. Stiller, C.Vögtli, B. Plattner: *Reservation-based Charging in an Integrated Services Network*; 4th INFORMS TELECOMMUNICATIONS Conference, Boca Raton, Florida, U.S.A., March 11-18, 1998, Session MC-2.
 - [9] D. Ferrari, L. Delgrossi: *Charging for QoS*; 6th IEEE/IFIP International Workshop on Quality of Service, Napa, California, U.S.A., May 18–20, 1998, pp. vii–xiii.
 - [10] A. Gupta, D. O. Stahl, A. B. Whinston: *Managing the Internet as an Economic System*; CISM, University of Texas at Austin, U.S.A., July 1994.
 - [11] S. Herzog: *RSVP Extensions for Policy Control*; Internet Draft, Work in Progress, January 1999.
 - [12] Hewlett-Packard: *HP Smart Internet Homepage*; <http://www.hp.com/smartinternet>, 1994-2000.
 - [13] Hewlett-Packard: *SIU Documentation*; 1999.
 - [14] ISO/IEC: *Guidelines for the management of IT Security (GMITS) - Part 1: Concepts and Models for IT Security*, ISO/IEC TR 13335-1, 1996.
 - [15] ISO/IEC: *Guidelines for the management of IT Security (GMITS) - Part 2: Managing and Planning IT Security*, ISO/IEC TR 13335-2, 1997.
 - [16] ISO/IEC: *Guidelines for the management of IT Security (GMITS) - Part 3: Techniques for the Management of IT Security*, ISO/IEC TR 13335-3, 1998.
 - [17] ITU-T Recommendation E.800: *Quality-of-Service and Dependent Vocabulary*; Geneva, Switzerland, November 22, 1994.
 - [18] ITU-T D-Series Recommendations: *General Tariff Principles - Definitions*; D.000, Geneva, Switzerland, 03/93.
 - [19] M3I Project Consortium: *M3I Annex I - Description of Work*; October 13, 1999.
 - [20] J. MacKie-Mason, H. Varian: *Pricing Congestible Network Resources*; IEEE Journal on Selected Areas in Communications, Vol. 13, No. 7, 1995, pp 1141 – 1149.
 - [21] M. Karsten (Ed.): *Pricing Mechanisms Design (PM)*; M3I Deliverable 3, Version 1.0, June 30, 2000.
 - [22] M. Karsten, J. Schmitt, B. Stiller, L. Wolf: *Charging for Packet-switched Network Communications - Motivation and Overview*; Computer Communications, Vol. 23, No. 3, February 1, 2000, pp. 290-302.
 - [23] M. Karsten, J. Schmitt, L. Wolf, R. Steinmetz: *An Embedded Charging Approach for RSVP*; 6th IEEE/IFIP International Workshop on Quality of Service, Napa, California, U.S.A., May 18–20, 1998.
 - [24] S. Kent, R. Atkinson: *Security Architecture for the Internet Protocol (IPSec)*; RFC 2401, Request For Comments, Standards Track, November 1998.
 - [25] H. Kneer, U. Zurfluh, G. Dermler, B. Stiller: *A Business Model for Charging and Accounting Internet Services*; International Conference on Electronic Commerce and Web-technologies (EC-WEB 2000); Greenwich, England, U.K., September 4-6, 2000.

- [26] D. Morris, V. Pronk: *Charging for ATM Services*; IEEE Communications Magazine, Vol. 37, No. 5, May 1999, pp 133-139.
- [27] D. Ranasinghe, J. Norgaard: *Charging Generalized: A Generic Role Model for Charging and Billing Services in Telecommunications*; in "Intelligent networks", Edt. J. Harju, T. Karttunen, O. Martikainen, Chapman Hall, London, England, 1995, pp. 159 – 172. Also in: IFIP Workshop on Intelligent Networks, Lappeenranta, Finland, August 8-9, 1994.
- [28] B. Schneier: *Applied Cryptography*; Wiley and Sons, New York, U.S.A., 1994.
- [29] S. Shenker, D. Clark, D. Estrin, S. Herzog. *Pricing in Computer Networks: Reshaping the Research Agenda*; ACM Computer Communication Review, Vol. 26, No. 2, April 1996, pp. 19–43.
- [30] D. Songhurst (Edt.): *Charging Communication Networks - From Theory to Practice*; Elsevier Publisher, Amsterdam, The Netherlands, 1999.
- [31] B. Stiller: *Pricing and Charging of Packet-based Networks*; Tutorial T12 presented at IEEE/IFIP Network Operations and Management Symposium (NOMS 2000), Honolulu, Hawaii, April 14, 2000.
- [32] B. Stiller, T. Braun, M. Günter, B. Plattner: *The CATI Project: Charging and Accounting Technology for the Internet*; 5th European Conference on Multimedia Applications, Services, and Techniques (ECMAST'99), Madrid, Spain, May 26-28, 1999. LNCS, Springer Verlag, Heidelberg, Germany, Vol. 1629, pp. 281-296l.
- [33] B. Stiller, G. Fankhauser, N. Weiler, B. Plattner: *Charging and Accounting for Integrated Internet Services - State of the Art, Problems, and Trends*; The Internet Summit (INET '98), Geneva, Switzerland, July 21-24, 1998, Session Commerce and Finance, Track 3.
- [34] B. Stiller, P. Reichl (Edt.): *Cost Model*; M3I Deliverable 8; Version 0.4, March 15, 2000.
- [35] M. Webster: *Webster's Collegiate Dictionary*; 10th Edition, Springfield, Massachusetts, U.S.A., 1996.
- [36] R. Yavatkar, D. Pendarakis, R. Guerin: *A Framework for Policy-based Admission Control*; RFC 2753, Informational Request For Comments, January 2000.

12 Abbreviations

A	Accounting Component
AH	Authentication Header
ATM	Asynchronous Transfer Mode
B	Billing Component
BT	British Telecommunications Research, Ipswich
C	Charge Calculation Component
CAS	Charging and Accounting System
CATI	Charging and Accounting Technology for the Internet
COPS	Common Open Policy Service
DB	Database
CPN	Customer Premises Network
EPC	Enterprise Policy Control
ESP	Encrypted Secure Payload

ETHZ	Eidgenössische Technische Hochschule Zürich, Switzerland
FSM	Finite State Machine
HP	Hewlett-Packard European Laboratories, Bristol
IETF	Internet Engineering Task Force
ID	Identification
IN	Intelligent Networks
IP	Internet Protocol
IPSec	Internet Protocol Security Architecture (Protocol)
ISP	Internet Service Provider
M	Metering component
M3I	Market Managed Multi-service Internet
N-ISDN	Narrowband Integrated Services Digital Network
NRD	Non-repudiation of Delivery
NRO	Non-repudiation of Origin
PDU	Protocol Data Units
PIN	Personal Identification Number
QoS	Quality-of-Service
R	Router
RAID	Redundant Array of Inexpensive Disks
RAP	Resource Allocation Protocol
RFC	Request for Comments
RSVP	Resource Reservation Protocol
SIU	Smart Internet Usage
SLA	Service Level Agreement
TCP	Transmission Control Protocol
ToS	Type-of-Service
TUD	Technische Hochschule Darmstadt, Germany

13 Acknowledgements

The overall design architecture for the CAS Design (WP6) as well as the Pricing Tools Design (WP5) has been developed by TUD and ETHZ based on most basic components of the M3I architecture description language (WP3) provided by BT.

Therefore, subsections of this document have been provided by Martin Karsten, Jens Schmitt, and Oliver Heckmann from TUD. The main parts of this document have been developed and written by Jan Gerke, Placi Flury, and Burkhard Stiller and include various comments received from Peter Reichl (ETHZ). Bob Briscoe, Pete Bagnall, David Freestone, Ben Strulo, Trevor Burbridge, Kostas Damianakis, and Jerome Tassel from BT are acknowledged greatly for their subsection on mass customization. Finally, Huw Oliver and Hans Daanen from HP provided various helpful comments on the structure of the introduction as well as on the applicability of SIU for the CAS.