

A Customer Service Management Architecture for the Internet

Ron A.M. Sprenkels, Aiko Pras, Bert-Jan van Beijnum, and Leo de Goede

University of Twente
Computer Science department
P.O. Box 217, 7500 AE Enschede, The Netherlands
{sprenkel,pras,beijnum,goede}@cs.utwente.nl

Abstract. Managing services on the Internet is becoming more and more complex and time consuming for service providers since services are increasing both in number and complexity. Also the number of users per service is going up. A solution to this problem is to allow the service users themselves to partly manage the services they are using. This is called Customer Service Management, and it will both increase the perceived value of the services to the users as well as lower the operational costs of service management for the service provider. This paper presents an architecture for Customer Service Management in the Internet.

1 Introduction

The Internet has evolved from a small network for universities and research organisations into the world-spanning network with many millions of users it is today. Not only the size of the Internet has changed but also the services that are delivered over it have changed. Transport services for instance are developing from straight forward best effort packet forwarding services into services that support different Quality of Service classes. An example of this are differentiated services currently under development in the diffserv working group of the IETF [4]. The increasing complexity of services gives rise to more service parameters that can be configured differently per service user.

As a result of this development managing the Internet and the services supported by it has become more difficult and in particular more labour intensive for the service provider organisations. Not only the number of subscribers to a service is increasing, but also the amount of configuration information per subscriber. This information might need to be changed relatively frequently.

To address this service management issue, this paper proposes a Customer Service Management Architecture (CSMA), based on work done earlier in [5]. Customer Service Management (CSM) transfers a part of the task of managing a service from the service provider to the service customer. The goal here is to reduce the amount of service management activities the service provider has to perform (thus lowering the operational costs) and at the same time to increase the value of the service to the service subscriber.

This paper is organized as follows. Section 2 provides an introduction to Customer Service Management. The terms Service Level Agreement and Service Level Specification are discussed, and the requirements for a Customer Service Management Architecture are presented. Section 3 presents the Customer Service Management Architecture, and discusses its components in detail. Section 4 gives a brief overview of related work that has been done in this field. Section 5 presents some conclusions and possible directions for future work.

2 Customer Service Management

Figure 1 shows an overview of the relevant components for Customer Service Management.

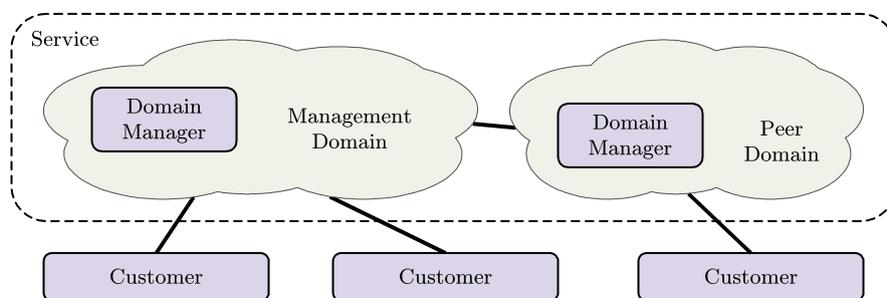


Fig. 1. Customer Service Management Components

A Service Provider organisation is an autonomous management domain, and it delivers a service to its customers. The responsibility for the management of the domain resides with the Domain Manager. From a management perspective the Internet can be seen as a interconnected set of such autonomous domains. To deliver a single service to a set of customers of different service providers, multiple interconnected service providers will cooperate. These service providers are assumed to be peer domains of each other. One such peer domain is shown in the figure, with its associated domain manager. The dashed line represents the service boundary between the set of management domains providing the service and the set of customers using the service.

To illustrate both the requirements and the architecture for Customer Service Management we will be using an example service: the Gold Silver Bronze IP Transport Service (or GSB service for short). It is a basic IP connectivity service that supports three different QoS levels. We define the GSB service such that a customer of this service can send and receive IP traffic in only one of the three levels at a given time. The service has a bandwidth parameter that specifies the bandwidth (number of bits per second) configured for that class. The classes have different characteristics for delay and loss probability. Gold category is a

low loss probability, low delay service. Silver category is a low loss probability, higher delay and delay variation service. Bronze category is a best effort service, so it can have both higher delays and a higher loss probability.

2.1 Service Level Agreement

The Service Level Agreement (SLA) contains all of the agreements between the service provider organisation and the customer about the delivery of the service [12]. The SLA can be considered to consist of a static and a dynamic section.

The static section contains all the legal arrangements for the service subscription. An example arrangement is that the service provider guarantees that the service will be available 99.9 percent of the time per month, and that a service outage will not last for longer than one hour. The SLA could also specify what happens when the service provider does not meet this guarantee, like paying back a part of the service subscription fee for instance.

The static section also contains the boundaries and the possible options for all service parameters that the customer can set. The SLA for the GSB example service contains for a user what the minimum and the maximum bandwidth parameter values are for each QoS class. From a technical point of view, users of the GSB service can send and receive traffic in three different QoS classes. What kind of QoS classes the service user can actually use is however not only a technical issue but also to a large extent an issue of business strategies and marketing. The service provider organisation can make several different service offerings that all map onto the same technical GSB service infrastructure. To illustrate this we give two example variants of the GSB service.

The "GSB business" service subscription puts the service user in the Silver class by default and it includes 10 hours per month of Gold service for no additional charge. Additional time in gold service costs two dollar per hour.

The "GSB economy" service subscription puts the service user in the Bronze class by default, it includes 5 hours of Silver class per month for no additional charge, and one dollar per hour for additional hours. Gold class is available for five dollar per hour.

As the names suggest, the business variant will have a higher monthly standard charge than the economy variant. The monthly bill for both service variants is usage sensitive; the customer pays a monthly subscription fee and an additional fee per unit time spent in higher than standard service categories. The SLA also contains provisions for how payments are made, how and where the service subscriber can complain if the service provider does not meet the minimum quality levels agreed upon in the SLA, and so on.

2.2 Service Level Specification

The SLS is the dynamic part of the Service Level Agreement between the service provider and the customer. The SLS has the current values of every customer configurable parameter. These parameters specify the more technical oriented

properties of the service, like for a transport service the quality of service class, the available bandwidth in that class etc. The SLS is dynamic because the values specified in it can change on a much shorter time scale than the agreements in the static part of the SLA. When service users have the ability to easily change the settings in their SLS, changes might occur on a time scale measured in minutes. It might even be required to take measures to prevent too frequent changes in the SLS parameters, to prevent customers from overloading the mechanism that is used for changing those parameters. This can be achieved in different ways, for instance by specifying and enforcing a minimum time between two SLS parameter changes, or by charging a small fee every time SLS parameters are changed.

For the GSB example service, the SLS contains which of the Gold, Silver or Bronze service classes the user currently has selected, and how much bandwidth within that service class is selected.

2.3 CSM Requirements

The key idea behind going from Service Management to Customer Service Management is to delegate a part of the task of service management from the service provider organisation to the service customer. For this purpose the SLA is extended with provisions that determine what kind of service management actions the customer can do, and what the boundaries of any parameters to such actions are. This gives rise to the following requirements for a Customer Service Management Architecture:

1. The SLA and SLS determine the current configuration of the service for a particular customer.
2. The Domain Manager is responsible for the management of the domain, and can create and change SLAs for this purpose. The CSM architecture therefore needs to provide the domain manager with the possibility to fully manage the SLAs of all the customers in the domain.
3. The SLA specifies what the customer can do with his SLS. The CSM architecture needs to provide every customer with the possibility to fully manage his SLS, while at the same time enforcing the constraints of the SLA.
4. For every customer the service has to behave as specified in that customer's SLA and SLS. The CSM architecture has to provide the necessary mechanisms to map the SLA and SLS settings onto the actual configuration of the service provider equipment.
5. Service customers will want information that allows them to determine if they are getting the service levels they are paying for. The CSM architecture needs to provide customers with this information.

3 Related Work

This section outlines two of the areas where work related to the Customer Service Management Architecture presented in this paper has been done. The areas

of interest are the International Telecommunication Union (ITU-T) and the Internet Engineering Task Force (IETF). Where appropriate it will be indicated why the requirements for Customer Service Management as outlined in section 2.3 are not fully met.

3.1 ITU-T

The Telecommunication Standardization Sector of the International Telecommunication Union (ITU-T) in recommendation M.3010 [11,10] outlined the principles for a Telecommunications Management Network (TMN). Although the TMN set of recommendations are mainly targeted at telephone type networks, the underlying principles of customer service management can be applied to Internet management as well.

Recommendation M.3010 describes (apart from the general description of a single TMN) in section 2.6 how multiple TMNs can access each other, and how network users can access the TMN functions of their network. The latter case is called Customer Network Management (CNM), and is described in four separate recommendations. In [8] the architecture for Customer Service Management is specified. In [9] the CNM service definitions are specified; these are organized into six functional areas: the five areas known as the FCAPS functional areas (Fault, Configuration, Accounting, Performance and Security), and the CNM supporting services area. For delivering the CNM services there are currently interface definitions for CMIP and for EDI, defined in Recommendations X.161 and X.162.

The TMN recommendations specify an X reference point and associated X interface to allow an Operations Systems Function (OSF) from one TMN to access an OSF in a different TMN. Figure 2 shows how the TMN of a customer interacts over a X interface with the TMN of a service provider.

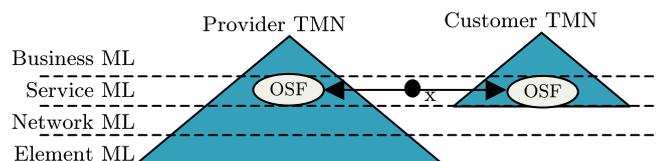


Fig. 2. Customer Service Management in the TMN Context

The interaction between the Customer TMN and the Provider TMN occurs between the Operations System Functions (OSFs) contained within the Service Management Layer at each side of the X reference point. The level of control that the provider gives to the customer is determined by the actual implementation of the X reference point for a particular service.

From this overview it can be concluded that the TMN management model supports the notion of Customer Service Management as defined in this paper.

However TMN CSM as defined in the current the ITU-T documents can not immediately be used in an IP-Internet environment, for a number of reasons. First, the service definitions are tailored to the ITU-T notion of data networks, not to IP networks. This means that the services are targeted on connection oriented ISDN data networks, use telephone numbers for addressing and so on. It is in general not straight forward to map these services onto their possible IP network counterparts. In IP environments, CSM functions are needed that allow customers to manage their diffserv subscription, their web hosting service subscription, the GSB example service introduced in section 2, and so on.

Second, quite a large portion of the services mentioned in [9] are not described in detail, but are marked as ‘for further study’, so they cannot be used as a starting point to define IP network CSM services.

Third, protocol mappings for these services are currently defined for CMIP [6] and for Electronic Data Interchange (EDI) [7]. These protocols are not widely supported by IP equipment.

3.2 The Internet Engineering Task Force

In the Internet Engineering Task Force (IETF) work is being done that is closely related to service management and customer service management. In September 1999 the IETF chartered a design team to identify a global set of configuration management requirements for IP-based networks. The findings of this team resulted in the creation of the “Configuration Management with SNMP” working group [3] in the Operations and Management area. This working group is examining how SNMP [1] can be used for configuration management. The model as proposed the working group for configuring devices in an IP network is given in Figure 3.

The purpose of the model is to convert from “Configuration Management Data (High-level Policies)” down to device local configuration for every device. This is done via an intermediate set of data, the “Network-wide Configuration Data”. To create this data set also network topology and network status and performance information is used. The final step translates the intermediate data into device local configurations by means of an “Configuration Data Translator”. The combined procedure therefore covers one of the two primary tasks of the CSMM proposed in this paper, which is to convert joint SLA/SLS information into device configurations. The other primary task, which is to map device feedback information to the relevant SLA/SLSs is not directly covered here.

The Differentiated Services (diffserv) working group is another IETF working group that does work related to service management. The working group defined a SNMP Management Information Base (MIB) to allow Policy Based management of a diffserv capable router [2]. In this MIB Policy Rule Classes and Policy Rule Instances are used to describe the intended configuration of the router, at a higher conceptual level than that of individual interfaces of individual diffserv routers. The mapping that occurs in the router from Policy Rule Instances to the actual configuration of a interface is similar to the mapping that occurs in the CSMM presented in this paper.

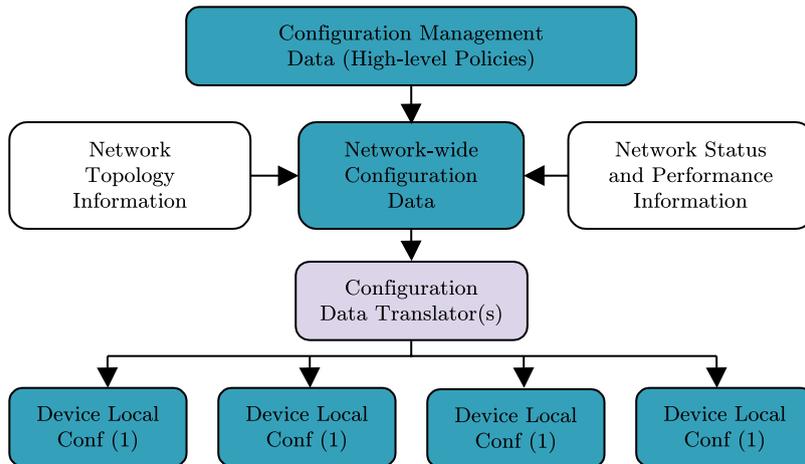


Fig. 3. IETF Configuration Management Model

With respect to the CSM requirements (section 2.3) the following about the IETF work can be observed. The configuration management model enables the domain manager to specify the target behaviour of services provided by the management domain. Depending on the nature of the high level policies, this covers requirements 1 and 2. The configuration data translator eventually maps the high level policies onto the configuration of individual devices, this covers requirement 4. The requirements 3 and 5 are not addressed by the IETF work: the configuration management model does not have provisions that allow individual customers to manage their own SLS settings. Also customers do not receive any feedback on whether the service is meeting the requirements specified in the SLA/SLS.

4 The Customer Service Management Architecture

Figure 4 shows an overview of the Customer Service Management (CSM) Architecture. The central component of the CSM architecture is the Customer Service Management Module (CSMM). Each domain has its own CSMM. The Domain Manager, the Customers of the domain and Peer Domain Managers (these are domain managers of adjacent peer domains) all access this module to do their respective share of managing the domain. This is depicted by the thick solid arrows. The CSMM in turn configures and controls the equipment (like routers, servers etc.) of the domain, and it receives feedback information from that equipment. This is also represented by thick solid arrows. For IP connectivity, the customers are of course also connected to network elements inside the management domain directly, and the network elements are connected with each other. This is represented by thin dashed lines.

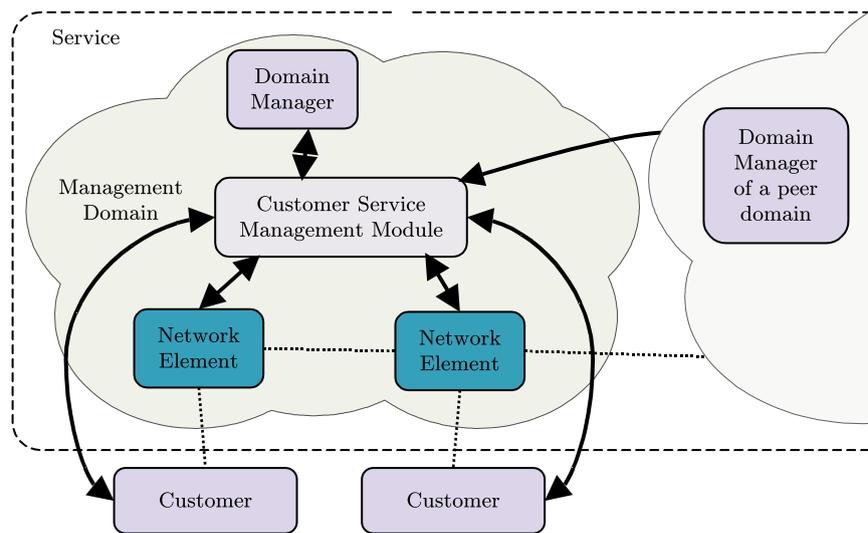


Fig. 4. Overview of the Customer Service Management Architecture

4.1 Customer Service Management Module

The Customer Service Management Module (CSMM) (Figure 5) is responsible for maintaining the domain configuration. Its operation revolves around two sets of information (modelled as data bases), the SLAs and the SLSs. They contain the SLA and SLS settings of every domain customer. The CSMM allows domain managers to create, modify and delete SLAs, and it allows customers and peer domain managers to modify their SLS settings. The CSMM also configures the domain equipment based on the current SLS values, and it receives feedback information from the equipment to relate back to the customer SLAs and SLSs.

The domain managers access to the database of SLAs is subject to a checking function. The CSMM enforces the SLA values to be consistent with what the service can actually provide. For example, it checks if the bandwidth parameter of the example GSB transport service for a customer is not higher than the access speed of that customers access link. This check requires detailed knowledge of capabilities of the service providers equipment and its internal structure. Also the resources already set aside for other SLA guarantees must be know to be able to perform this check. More on the domain manager in section 4.2.

The customer can read and modify the parameter values of its own SLS. This access is also subject to a checking function; the CSMM checks if the requested new SLS values are within the boundaries of the values specified in the SLA. The customers access to the CSMM is discussed further in section 4.3.

The peer domain managers are similar to a customer; a peer domain has a SLA with this domain, and the peer domain manager can modify its SLS. A

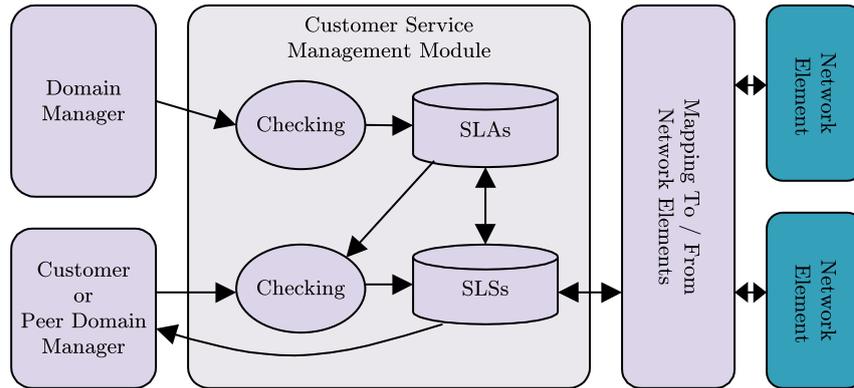


Fig. 5. Internal Structure of the Customer Service Management Module

peer domain SLA will differ from a customer SLA. Also a peer domain will in general be connected to the domain differently, with a high capacity physical connection for instance. More on the peer domain in section 4.4.

The final step in making sure that the service is delivered in accordance with the provisions in the SLA and SLS is to configure the service provider equipment in such a way that these provisions are met. This requires a mapping from the joint SLS values onto the equipment configuration (section 4.5).

4.2 Domain Manager

The domain manager is overall responsible for the operation of the domain, and therefore sets the SLA parameters for the customers and for the connected peer domains. A SLA is created upon a customer subscribing to a service is changed only infrequently, if at all. Information not only flows from the domain manager to the CSMM, but also in the other direction. The CSMM knows at any point in time what SLS service parameters are in effect for a particular customer, because the only way to modify these parameters is through the CSMM. The domain manager can use this information for instance for accounting purposes.

It is to be expected that the SLA information that is kept in the CSMM has a close relationship with other information processing and administrative systems outside of the CSMM. It is therefore not likely that the domain manager will only use a graphical user interface; instead some form of automated information maintenance interface to the SLAs contained in the CSMM will be needed.

4.3 Customer

The customer can adjust his SLS parameters within the boundaries given in the SLA. This type of change occurs more frequently than SLA changes. To change SLS parameters the customer needs to access the parameter setting function block in the CSM module. There are different options for this discussed next.

One option for the communication between the customer and the CSMM is to use standard web technology. Web technology in this context means that the customer is a human using a web-browser to access a web server inside the CSM module. Standard HTML web pages, HTML fill-out forms and java applets can be used. Essentially this means that a web browser is used to implement the client in a client server system, where the server is contained in the CSMM.

A benefit of using a standard web browser to do CSM is that web browsers are commonly available today. In particular when customers are private persons that are also the only service user this can be a good solution, since it does not require installation of any new software on the customer side, and customers are already familiar web browsers. A down side might be that while the access mechanism is a standard one, the user interface to the SLS configuration is not. Different implementors are free to design their own application for SLS management, and as a result two interface implementations to the same SLS parameters will differ. This will be confusing to customers that change from one service provider to another. Another down side is that a web based GUI is probably not convenient for automated SLS management.

Another option is to define a new protocol. This has the advantage that it can be tailored for the purpose, and that it can be designed to be also usable to do CSM in an automated way. A consequence of a new protocol is that new applications need to be developed and installed on the customers premises. Note that the customer can also have read-only access to the SLA.

4.4 Peer Domain Manager

The SLS management requirements for a peer domain differ from the SLS management requirements of a customer in a number of ways. First the SLA and SLS with a peer domain will be different from those with a customer. As a result, different kinds of parameters need to be managed there. Second, the frequency at which a peer domain changes its SLS parameters is probably higher. Third, the peer domain manager will in general be a piece of software instead of a human.

For these reasons the management of the SLS parameters of a peer domain will require a protocol type solution and not just a graphical user interface to a management application inside the CSMM.

4.5 Mapping Function

On one side of the mapping function there are the SLAs and SLSs of all the service users; on the other the devices that together construct the service provider. The CSMM mapping function translates information in both directions. It maps from SLA/SLS values to device specific configuration information, and from device specific feedback information to SLA and SLS specific feedback information in the device independent internal format of the CSM module.

The SLA and SLS information is available on a per-service-user basis. The service provider consists ultimately of network elements that need to be config-

ured. The mapping function therefore has to determine on the basis of the joint SLSs of all users how to configure all of the domain network elements (Figure 6).

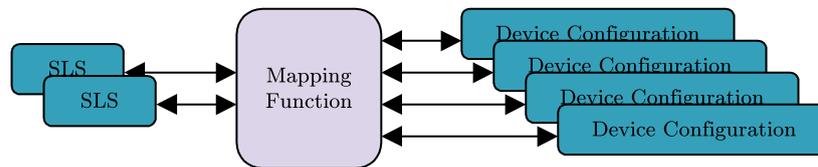


Fig. 6. Mapping from Multiple SLSs to Multiple Device Configurations

The configuration of a specific device depends on the SLS parameters of a number of service users. The mapping function therefore creates and maintains the configuration of each device based on the SLSs of multiple service users.

Service users get feedback information from the service provider for instance about how the actual achieved QoS values compare to the ones specified in the SLA. The feedback information of a single device will in general be relevant to multiple service users and their SLAs. In terms of the model of a domain this is a large distance to cover; in a way it is a form of mapping service management requirements directly on element management primitives.

The service provider can consist of devices of different vendors that do not have standardized interfaces for configuration. The mapping function maps from the CSM internal information to the format supported by a particular device. Likewise, the device feedback information can differ in semantics and format between vendors. The mapping function accommodates these differences, thus supporting service providers utilizing a mixture of devices from different vendors.

Detailed knowledge of the network configuration, topology, and load situation is needed to be able to map back and forth between per service user SLA/SLS information and network element information. It is expected that the mapping procedure will involve an intermediate stage in the form of a network management layer. The mapping function is therefore probably the most difficult part of the CSM Architecture that needs to be further investigated and developed.

5 Conclusions and Future Work

Customer Service Management allows the partial transfer of the task of service management from the service provider to the service users. The Service Level Agreement and Service Level Specification play a central role in CSM since they determine in detail the possible user service management actions. This paper presented an architecture for Customer Service Management for the Internet. Application of this architecture to service management has two main advantages.

First, the service providers can achieve lower costs of operating their services. Because a portion of all service management actions is now done by the service

users themselves, without human interference of the service provider, the service provider saves on labour costs.

Second, when using CSM the service user can control his own service more closely. He can make changes to the service configuration at any time of day, because at the service provider no human interaction needs to be involved. It is expected that the increased level of control results in a higher perceived value of the service to the service user.

In the ITU-T work has been done on CSM, but since this work is targeted on connection oriented ISDN data networks, the CSM functions defined there are not directly usable for IP networks.

The IETF work on configuration management allows a domain manager to configure all network elements in his domain based on a set of high level policies, but not individual customers making changes to their SLA or SLS settings.

The central component of the CSM architecture presented in this paper is the Customer Service Management Module (CSMM). This architecture is still at a fairly high level. In particular, important details about the mapping function between service user SLA/SLS information and network elements need to be worked out. An example like the GSB service (section 2.1) needs to be worked out in detail, for a clearly defined service, a set of service users and a network topology. Such an example should then be implemented to gain practical experience and to further validate the architecture and its underlying assumptions.

References

1. Jeffrey Case, Russ Mundy, David Partain, and Bob Stewart. Introduction to version 3 of the Internet-standard network management framework. RFC 2570, April 1999.
2. Michael Fine et. al. Differentiated services quality of service policy information base. internet draft, draft-ietf-diffserv-pib-01, work in progress, July 2000.
3. IETF. Configuration management with SNMP. <http://www.ietf.org/>, 2000.
4. IETF. Differentiated services. <http://www.ietf.org/index.html>, 2000.
5. ING. ING WU2 new Internet management architectures. <http://ing.ctit.utwente.nl/WU2>, July 2000.
6. ITU-T. Common management information protocol specification for CCITT applications. Recommendation X.711, 1991.
7. ITU-T. Message handling: Electronic data interchange messaging service. Recommendation F.435, 1991.
8. ITU-T. Architecture for customer network management service for public data networks. Recommendation X.160, July 1994.
9. ITU-T. Definition of customer network management services for public data networks. Recommendation X.161, July 1994.
10. ITU-T. Principles for a telecommunications management network. Recommendation M.3010, May 1996.
11. Aiko Pras, Bert-Jan van Beijnum, and Ron Sprenkels. Introduction to TMN. Technical Report No. 99-09, Center for Telematics and Information Technology, P.O. Box 217, 7500 AE Enschede, The Netherlands, April 1999.
12. Dinesh Verma. *Supporting Service Level Agreements on IP Networks*. Macmillan Technical Publishing, 201 West 103rs Street, Indianapolis, IN 46290 USA, 1999.