Luciano Gavi

Service Level Management for IP Networks

This paper gives a general overview of the service level management (SLM) approach, which can offer significant advantage to service providers competing in the current European market. Starting from a theoretical approach, concepts are then applied to telecommunications scenarios focusing on the provision of data services and the Internet protocol (IP). Major concepts related to service level agreements (SLAs) and SLA management systems are analysed in greater detail defining how such concepts can be applied to IP networks. Finally some details are given concerning possible parameters to be applied to the provision of IP services, including both permanent and switched services.

Introduction

In the current European telecommunications environment the advent of a liberalised market gives rise to complex scenarios where a number of players are involved in providing services to end-customers. Players include network carriers ensuring the availability of infrastructure, network operators offering basic transport services, service providers providing basic and/or advanced services, service brokers etc.

The study of the complicated interactions that may occur among different players in providing telecommunications services to end-customers is outside the scope of this paper.

Luciano Gavi:
Systems Division,
Sirti S.p.A.
Cassina De' Pecchi
Milano, Italy
Tel: +39 02 6677 3699
Fax: +39 02 6677 3215
E-mail: L.Gavi@sirti.it

What is worthwhile mentioning here is that organisations tend to categorise their telecommunications assets and activities in terms of services that are supported and relevant service levels that can be monitored. Analysis of service levels means managing, in a business oriented perspective, the network technologies and applications adopted as well as the administrative and organisational structure that enable service life cycles.

This paper starts by describing some fundamental concepts of service level management. Then such concepts are applied to the telecommunications environment, focusing on data networks. The analysis is not intended to be exhaustive but the paper tries to give some flavour of the importance that SLM is getting for data networks. Some details are given about service level agreements and related management systems that will soon represent key issues in the support of IP services.

Basic Concepts of Service Level Management

The aim of this section is to describe the SLM approach and related concepts. Performing SLM means describing the activities carried out in terms of services made available to some specific customers. For each service some parameters are identified that allow the definition of metrics and service levels. Thus the monitoring of the parameters associated to a service gives a measure of the performances related to the service being provided.

This approach focused on SLM gives rise to substantial benefits in terms of operational efficiency and effectiveness, allowing a continuous comparison of service levels achieved against levels targeted or requested by customers.

SLM can be applied to different environments. For instance the provision of IT services within an organisation can be managed using SLM. SLM can be applied also to other fields like transportation and energy. In the following more details are given about SLM in general. Then SLM concepts are applied to the telecommunications field.

Defining service level management

To implement SLM the following aspects should be taken into account:

- definition of service level components or *objects* that can represent the service with a number of parameters showing how well the service is provided;
- definition of the SLM policy to be adopted determining rules, activities and constraints;
- identification of a proper organisation ensuring that the right people are involved at the right time to carry out SLM;
- definition of formal service level agreements; and
- identification of proper operations support systems, if any.

The definition of SLM objects can vary depending on the environment involved but the related parameters usually describe such aspects as availability and quality of service. Aspects related to the object's criticality and time-of-day relevance are also taken into account. Objects and related parameters have to be instantiated to represent real service instances.

As an example in the IT environment an SLM object can be a workstation and related parameters can be system availability and file system status.

The definition of policy comprises the specification of the processes and interactions that should run to ensure SLM operations. Processes include:

- monitoring, which describes how the SLM objects' parameters are measured and collected;
- data processing, which describes how data made available by the previous process are aggregated and elaborated;
- services maintenance, which describes which actions should be carried out (acting either proactively or in reaction to problems) to guarantee adequate and constant levels of services; and
- reporting, which describes the scope and format of data to be recorded and made available.

To enable SLM an appropriate staff of people has to be assembled, defining who is responsible for doing what and identifying a work flow that minimises possible arbitrary choices. Organising a proper staff is a hard job and can impact significantly on traditional structures, introducing new interactions and responsibilities.

A service level agreement (SLA) is a sort of contract that formally defines what the service provider commits to ensure to the customer in terms of SLM objects and related parameters. An SLA usually comprises some major SLM objects and related parameters that represent the service being provided. Considering the same example mentioned above, if a workstation is an identified SLM object, an SLA could be defined comprising that object and, as related parameters, system availability and file system status: such parameters can be measured and the following thresholds can be agreed upon:

- system availability > 99·8%
- file system occupancy < 90%

SLAs are usually associated with penalties that apply when agreed thresholds on service levels are exceeded.

Identifying clear SLA thresholds on well-defined objects and significant parameters enables a better relationship between a service provider and a customer:

- the service provider has to check continuously that thresholds are not exceeded (aiming if possible to keep some conservative distance from the thresholds); and
- the customer is aware of which level of services can be expected and is allowed to claim only if thresholds are exceeded.

To establish effective SLM, the adoption of an operations support system (OSS) can be very helpful. This becomes fundamental if SLAs are defined. Only the adoption of a proper OSS permits the service provider to monitor the service levels and adopt preventive and quick corrective actions to avoid penalties. Furthermore, use of SLA management OSSs enables the performance of the resources that are involved in the service provision to be studied, resources to be tuned to get maximum performances and differentiated levels of services for different customers to be dynamically defined.

Details about OSSs performing SLA management on telecommunications networks are given in a later section.

Service Level Management Scenarios in a Telecommunications Environment

All the SLM concepts illustrated above, including the definition of policies and the identification of a proper personnel structure, can be applied to a telecommunications environment. This paper, however, aims at focusing on the definition of SLAs and the architectural description of SLA management systems. Such aspects are becoming more and more important in the field of data networks because of the specific nature of those networks.

Basically, two main scenarios can be mentioned here that occur in the field of telecommunications networks:

(*a*) interactions occurring between an operator and the end customer; and
(*b*) interactions occurring between two operators (either network operators or service operators).

In case (*a*) the operator and the customer agree to define a formal contract referred to an SLA which clearly states all objects involved and expected levels of their parameters.

In case (*b*) the two operators interact in such a way that, at least for any specific exchange of data, one is acting as the service provider and the other as the customer. In this case the agreements that can be defined are referred to as *operational level agreeements* (OLAs). From a conceptual viewpoint an OLA is not so different from an SLA, but from a practical viewpoint some relevant differences may apply. Firstly an OLA is often tied to some extent to existing regulations and laws, while in the definition of an SLA a major role is played by commercial and marketing aspects. Furthermore, in several cases (for instance when an incumbent and a newcomer interact) service level objectives for the SLAs can be deeply affected by the service level objectives defined in OLAs.

For the sake of simplicity only case (*a*) will be considered in the following, focusing on an operator which plays the role of service provider (SP) providing services to end customers.

The adoption of SLAs has become consolidated when transmission services (for example, leased lines) are involved. In this case a few metrics can be taken into account—provisioning time, recovery time, availability time—and the models to define an SLA can be simple enough. More complex scenarios will occur when availability parameters calculated according to sophisticated algorithms (based on ITU M.2100 for transmission lines) are adopted. Similarly, when basic voice services are provided (for example, interconnection of new operators to the access network of an incumbent) some standard metrics can be defined that can easily be monitored by using capabilities available in voice switches of central offices.

Scenarios including the provision of data services are usually much more complicated. In this case the definition of SLM objects and the handling of related parameters can become difficult tasks.

A key issue is the collection and accounting of measures related to the objects' parameters. This can be accomplished in several ways:

- recording of rough data in major network components and batch transfer;
- collection of more sophisticated and aggregated data through intelligent devices; and
- use of distributed probes disseminated in strategic sites of the network

Such data can be collected taking into account that real time is not required but a reasonable time interval should be adopted if proper reactions to failures and degradation are addressed. Collection and accounting of measures data is a particularly complicated issue when IP services are involved. A specific analysis can be developed describing topics such as the identification of IP flows across a network and the implementation of application aware devices.

The need for deploying OSSs arises for data networks, as the complexity of interactions involved cannot be handled without such systems.

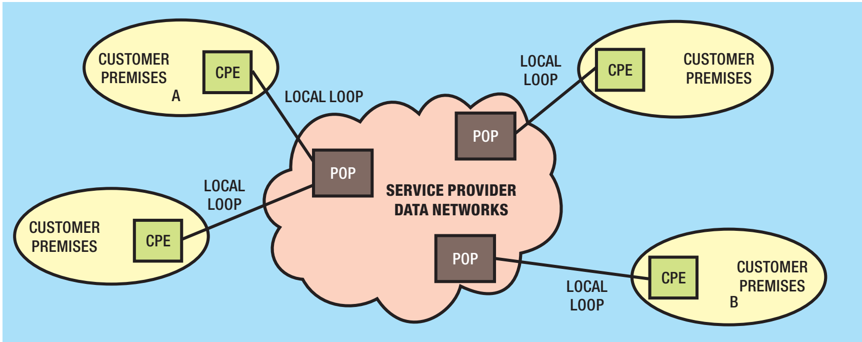

Figure 1—Architecture for data service provisioning

Modelling Service Levels Agreements

Defining an SLA involves a mixture of administrative and technical issues. Administrative issues refer to the definition of all terms and conditions that make the SLA contract applicable and enforceable including identification of validity times, responsibilities, exceptions, auditing policies, penalties, etc.

As far as technical aspects are concerned, a major issue is to define the right objects to represent a service for the purposes of the SLA. Such objects are also called *service access points* (SAPs). For instance, in the provision of frame relay or ATM services, permanent virtual circuits (PVCs) can be identified as the SAPs. Logical and/or physical local area network (LAN) interfaces can be considered when IP services are involved.

Two approaches are possible for defining and handling SAPs:

- identifying SAPs as logical entities only, related to the provision of a service; and
- identifying the mapping between a SAP as a logical entity and the underlying physical components that are involved in the provision of a service.

The approach adopted has relevant consequences on techniques that have to be adopted to collect and manage SLA data and on the general SLM policy that can be implemented.

Considering all the above aspects described above a typical SLA can be summarised as including:

- customer identification and related data,
- contract details,
- services profile and related data,
- SAPs involved and related parameters (for example, SAP weight),
- services instances involved,
- time intervals applicable and exceptions handled,
- levels of service and penalties defined depending on SAPs and service profiles adopted, and
- reports to be made available to the customer.

Focusing on data services, the following architecture can be defined to represent the provision of a service to a customer located in sites A and B (see Figure 1).

A point of presence (POP) is defined as the interface between the local loop and service provider's network. In the case of IP services a POP can be an edge router for permanent services and a network access server for switched services.

Two different situations can be envisaged to define SLA in the above scenario:

- the service provider (SP) is in charge of service support POP to POP; or
- the service provider is in charge of service support end-to-end.

The second case applies when the local loop and customer premises equipment (CPE) are under the responsibility of the SP. This usually happens when the CPE is provided and managed by the SP itself. SLA thresholds can vary remarkably from the first case.

In order to manage complex contracts, SLA management systems can be implemented, as described below.

SLA management systems

Systems to carry out SLA management are now becoming more common. Such systems implement the modelling concepts described above, adopting different solutions. A general architecture of an SLA management system is given in Figure 2.

The following considerations can be applied when choosing a SLA management product:

- models implemented in the system (availability of a well-defined object model, mapping between logical entities and physical entities);
- scope of the systems in terms of networks technologies that are supported (synchronous digital hierarchy (SDH), frame relay, ATM, IP…);
- availability of a database (external versus internal, standard versus proprietary, etc.);

Figure 2—Architecture of an SLA management system

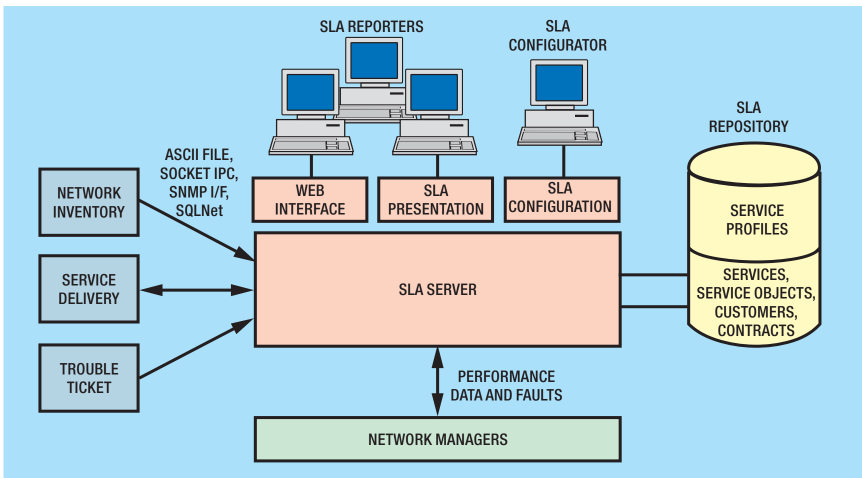

- configuration capabilities available (dynamic definition of SLAs' parameters, etc..);
- reporting capabilities available (easy definition of forms and reports, Web interface, etc.); and
- interfaces available to other OSSs (trouble ticket management, network inventory, service delivery, network performance management).

Effective interfaces with other OSSs like network inventory and service delivery are particularly important for IP services as in this case detailed information is to be exchanged concerning service/ customer features and SLA activation procedures.

If the system adopted is well focused on the services provided, significant advantages may occur including:

- automation of SLA administration activities;
- centralised control of the status of services being provided;
- full support for preventive and corrective maintenance with semi-real time monitoring;
- the possibility of defining SLAs differentiated for customers' location, time of day, etc.;
- recording of data on a dedicated repository and availability of sophisticated reporting; and
- the possibility of carrying out detailed analysis to tune network performances (capacity planning) and marketing strategies.

As far as IP networks, the adoption of a proper SLA management system, using the right models and sophisticated functions must be combined with appropriate data collection/accounting mechanisms. Implementing SLA management of IP VPNs and IP-based applications means managing high volumes of complex data.

SLA of IP Networks

The development of datacommunications networks is enabling the definition of sophisticated services for customers. Service have evolved from time-division multiplex (TDM) services through frame relay to ATM and IP services. Such technologies are much more complicated and require huge investments to ensure that networks are robust, flexible and manageable. Huge investments are however needed to compete in a liberalised market where customers can be captured only by a sound mixture of technological choices and marketing strategies.

As a result of the considerations developed in this paper investments must be applied not only to the network infrastructure but also to the definition of a proper service level management framework with a related SLA management system. This is the only approach that can guarantee a strategic advantage and a satisfactory control of cost to benefit trade-offs.

To give a flavour of the complex scenarios that have to be managed in providing IP services, a list of basic SLA parameters specific to IP services is given below (considering both permanent and switched services). Such parameters should be used in conjunction with the general parameters given in the section on 'Modelling Service Level Agreements'. Thus for instance proper measurement intervals should be identified as well as proper weights should be given to SAPs.

Parameters for permanent services:
- global availability of SAPs,
- SAP unitary availability,
- access guaranteed bandwidth,
- end-to-end guaranteed bandwidth,
- transfer delay, and
- IP packets loss/duplication rate.

Parameters for switched services:
- global availability of SAPs,
- SAP unitary availability,
- successful call rate,
- transfer delay, and
- IP packets loss/duplication rate.

In both cases, parameters related to maximum response time and maximum recovery time are usually considered as well.

Conclusions

Service level management is becoming more and more common as service providers are competing in the growing European data services market. A good approach to SLM implies the definition of an appropriate SLA and the implementation of a SLA management system. This is particularly important in the field of IP networks. However, the high level of complexity of such networks and the low operational expertise available to most service providers represent significant hurdles. Large effort is now devoted to clarifying the enabling technologies and the service scenarios that will rule the real advent of IP.

Having achieved a better understanding of the 'IP world', service providers will very soon need to differentiate from competitors and to pursue best control of their resources. At that stage definition of sound SLAs will become a must to attract customers. Implementation of SLA management systems will be necessary to manage a comprehensive portfolio of IP services.

References

All the concepts presented in this paper are the result of the work carried out by the author and his colleagues in the field of SLM over the last three years. Considerable experience has been achieved by performing requirements analysis for SLM OSSs to be either developed internally in Sirti or integrated taking third-party products. Technical documentation of all such systems (including applications for both IT and telecommunications environments) can be considered as a background of this paper. No reference to Sirti's and third party's products is given here because this is considered outside the scope of the paper.

Biography

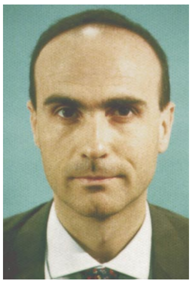

Luciano Gavi
Sirti S.p.A.

Luciano Gavi joined Sirti in 1989 and worked about five years in the research and development department, taking part in ESPRIT research projects devoted to the development of prototypes for new telecommunications networks (MANs). From 1992–1994, he participated in two European RACE projects, respectively devoted to telecommunications service engineering and security of telecommunications networks. In 1995, he joined Sirti Systems Division, participating, as a senior system engineer, in several bids in Italy and abroad in the field of network management. From 1996–1998, he managed projects devoted to the development of network management systems and ATM networking solutions. Since mid-1998, he has been a proposal manager, responsible for coordinating all projects related to network and service management for a major Sirti customer.