Comparative Study of Protocols for Dynamic Service Negotiation in Next Generation Internet

Venkatesh Sarangan[1], Jyh-Cheng Chen[2]

[1]Oklahoma State University, Stillwater, OK 74078, USA

[2]National Tsing Hua University, Hsinchu, Taiwan

Abstract

This paper presents a survey of the protocols that have been proposed for facilitating dynamic service negotiation in next generation Internet. We begin by illustrating the terms service level agreement (SLA) and service level specification (SLS) defined by the IETF. We then discuss the working of the existing service negotiation protocols with respect to a generic network architecture. Following that, we enumerate a list of characteristics desired in an ideal service negotiation protocol and draw a comparison between the various protocols based on this list. We conclude the paper by discussing possible future research directions in this area.

Key Words: Dynamic Service Negotiation, Next Generation Internet, Service Level Agreement, QoS.

Contact Author: Prof. Venkatesh Sarangan,

Computer Science Department, Oklahoma State University, Stillwater, OK 74078, USA.

Phone: +001-405-744-5672, Fax: +001-405-744-9097, Email: saranga@cs.okstate.edu

1 Introduction

The Internet is rapidly evolving into a commercial infrastructure, resulting in an increased demand for Quality of Service (QoS) guarantees from the network. To meet the service requirements of the next generation Internet (NGI) applications, new mechanisms for ensuring QoS need to be developed, requiring fundamental changes to the Internet's existing connectionless best-effort architecture. It is envisioned that in NGI, users will enjoy different levels of service for their traffic by executing contracts with their service providers. An important characteristic of NGI is to allow the users to dynamically adjust their desired service levels along with acceptable prices for the service. This feature is necessitated not only by the requirement to provide flexibility for the users, but also by the heterogeneity in the wireless sub-systems and end device capabilities.

A mobile user with ubiquitous connectivity is expected to have multiple radio interfaces, so that the user can choose the interface that provides the best instantaneous connectivity. While this will ensure ubiquitous connectivity, the mobile user will experience varied levels of service depending on the capacity of the wireless connectivity. It is also expected that a user will maintain connectivity through devices with diverse abilities. For example, a Personal Computer (PC) may be used at home or inside an office. While driving, a small handset will be more suitable and a Personal Digital Assistant (PDA) or laptop can be used when traveling. These devices differ not only in their processing and communication capabilities, but also in the applications they can run.

Thus, ubiquitous connectivity results in heterogeneous link layer technologies and diverse user terminals. In such a dynamic environment, it is not possible for a service provider to envision the service requirements of the users and provision the network accordingly. It is also difficult for the users to anticipate what level of service they really need. As the user will be charged for the services offered, a user would not want to pay for a high grade of service and not enjoy them due to limitations in the link layer or the device. Also, the user does not want to be serviced at a lower grade, if a higher grade of service is feasible. Such requirements can be satisfied only if the user is allowed to negotiate the service requirements dynamically.

Recognizing this need for dynamic service negotiation, several protocols have been developed by researchers towards this objective. This paper presents a survey of the various service negotiation protocols that have been proposed thus far, and presents a qualitative comparison of them. We highlight the similarities and differences among the protocols and also point out their strengths and weaknesses. Our paper is organized as follows. In section 2, we explain the term *service level agreement* and detail its essential attributes. In section 3, we describe the working of the existing service negotiation protocols. Section 4 elaborates on the characteristics of an ideal service negotiation protocol. Based on these features, section 5 provides a qualitative comparison of the various protocols. Finally, we conclude in section 6 with a discussion on the future research directions in this area.

2 Characterizing SLA and SLS

In general terms, service-level agreements (SLAs) are binding contracts between service providers and customers that define the services provided, the metrics associated with these services, acceptable and unacceptable service levels, liabilities on the part of the service provider and the customer, and actions to be taken in specific circumstances [1]. The IETF's (Internet Engineering Task Force's) DiffServ working group defines a SLA in networking parlance as follows: *A SLA is a service contract between a customer and a service provider that specifies the forwarding service a customer should receive* (IETF RFC 2475).

The SLA contains both technical and non-technical terms and conditions. The technical specification of the transport service is given in Service Level Specifications (SLSes). A SLS *is a set of parameters and their values, which together define the IP service offered to a traffic stream by a DiffServ domain* (IETF RFC 3260).

SLSes describe the traffic characteristics of IP flows and the nature of service to be offered by the network to these flows. SLSes could be of different types such as SLS for security services, SLS for mobility services, SLS for QoS guarantees, etc. As the focus of this paper is QoS, we concentrate only on SLS that are about QoS*. Primarily, such SLSes outline the QoS guarantees that are expected from a network for a given set of IP flows. In most (if not all) of the service negotiation protocols that have been proposed so far, it is this SLS that gets negotiated. Though the DiffServ working group defined SLS as a set of technical parameters, it did not specify these parameters. It was left to the discretion of individual service providers. However, based on the work that has been carried out so far, it is possible to consider a SLS as a composition of the following parameters [2,3]:

*Henceforth, the term SLS in this paper shall imply SLSes that are about QoS unless specified otherwise.

1. *Scope Specification:* The scope of a SLS refers to the topology or the geographical area over which the service has to be offered. The scope is specified by enumerating the ingress and egress routers for the traffic associated with the SLS. While scope specification helps the network to arrive at a better decision about its ability to satisfy a given SLA, it is not necessary that the SLA subscriber should be aware of the scope of his/her traffic. In case the subscriber does not have any knowledge about the ingress/egress points, he/she has the option of specifying a wild-card address for the ingress or egress points or both.

2. *Flow Specification:* A *flow* is defined as a distinguishable stream of related data packets requiring the same QoS. Flow specification identifies the flows that belong to the SLS under negotiation. The flows are identified by specifying values for one or more of parameters such as source and destination IP addresses, source and destination port numbers, and DiffServ code points (if applicable). Based on these specifications, the ingress router classifies the packets and marks them so that the packets belonging to these flows enjoy the negotiated service.

3. *Traffic Specification:* The traffic specification provides a description of the traffic volume of the flows associated with the SLS. Traffic flows could be characterized by different attributes such as peak rate, average rate, and minimum packet size. Based on these traffic specifications, the ingress nodes at the network do a conformance testing on the packets that belong to the specified flows. Only those packets that conform to the traffic specifications enjoy the negotiated service. Packets that do not conform to the negotiated traffic parameters can be dropped, shaped and/or remarked. A specific treatment for non-conforming packets can be also negotiated as a part of the SLS.

4. *Performance Guarantee Specification:* This describes the service guarantees that the flows identified by the flow ID and conforming to the traffic specifications will enjoy over the geographical region given by the scope. The service guarantees are given in terms of parameters such as delay, jitter, packet loss and throughput, and can be specified either *quantitatively* or *qualitatively*. A parameter is said to be quantitative if a definitive numerical value is assigned to it. A qualitative parameter is one whose value is set in comparison with others and cannot be quantified.

5. *Service Schedule Specification:* This elucidates the start and end times of the negotiated service, i.e.,

it indicates when the service will be available, e.g. 24×7, only during weekends, $9 - 5$ on weekdays, etc.

6. *SLS Identifier:* Each SLS may contain a unique identifier which helps in recognizing it. Though this parameter could be considered optional, including it offers flexibility in terms of enabling a client to re-negotiate a particular attribute of a SLS at a later time. The SLS ID may be decided by the service provider.

Apart from the above technical parameters that define the SLS, a SLA can also include several non-technical parameters as follows [4]:

1. *Price:* Price is an important parameter that needs to be negotiated. Note that as in other commodities, the price of a given network service can fluctuate over time and depends on network conditions. For example, under congested scenarios, the service provider can present a high charge for premium service to decrease the demand, while under light-loads, the same service could be offered for a discounted price to increase network utilization.

2. *Device and Network capabilities:* The SLA could contain a list of devices (such as mobile phones, PDAs, Laptop, PC, etc.) through which a subscriber is likely to interact with the network. This may allow the network to gauge the level of resources that the subscriber could probably negotiate, which in turn could allow the network to judge the admissibility of the SLAs of other subscribers in a better way. Similarly, the SLA might also have provisions for the subscriber to indicate the list of access technologies (such as GPRS, UMTS, WLAN, wired Ethernet, etc.) available to him/her.

3 Overview of service negotiation protocols

Having discussed the parameters that constitute a SLS and a SLA, we will now focus our attention on the protocols used for negotiating a SLS. The following protocols have been proposed by researchers for facilitating dynamic service[†] negotiation in NGI: (1) Resource Negotiation And Pricing (RNAP) protocol [5],

[†]By services, we refer to the types of SLSes that the network can support.

(2) Service Negotiation Protocol (SrNP) [6], (3) COPS-SLS [7], (4) Dynamic Service Negotiation Protocol (DSNP) [3], (5) QoS-NSIS‡ Signaling protocol (QoS-NSLP) [8], and (6) QoS Generic Signaling Layer Protocol (QoS-GSLP) [9]. Variants of some of the above protocols have also been proposed. While these protocols have quite a bit of commonality, they still have their own unique features. In what follows, we present a generic negotiation framework that all the protocols employ. Subsequently we discuss the features that are unique to each of these protocols.

3.1 A generic negotiation framework

The Internet is divided into different Autonomous Systems (ASes) or Domains administered by different service providers. Each domain possesses important network entities such as DHCP server, Authentication, Authorization, and Accounting (AAA) server, and Policy Information Base (PIB)§. In addition to the above, we assume that each domain includes two other entities namely, negotiation manager (NM) and resource manager (RM). NM is the entity that represents the domain in SLS negotiation with its subscribers. RM is the entity which decides on the admissibility of SLSes based on the resource availability in the domain. The various components present in the architecture are schematically shown in figure 1.

When a subscriber logs in to the network for the first time, he/she is notified by the NM about the pre-defined services available in the network. The notification can be either *pro-active* or *reactive*. The former refers to the case wherein the NM automatically informs the new subscriber about the services, while the latter refers to the scenario wherein NM sends the information only when the subscriber requests for the same. As soon as the subscriber obtains information about the available services, he/she can start the negotiation process with the NM. The subscriber can negotiate not only any of the pre-defined services, but also his/her own customized services. On receiving a service request message from the subscriber, the NM may consult with other entities such as RM and AAA server (if they are not co-located) to determine if the requested service can be provided. If so, the NM informs the appropriate network nodes, such as the edge routers (ERs), about the service to be rendered to the subscriber's traffic. The NM then sends a positive reply to the subscriber after which the subscriber starts enjoying the services. The NM may also contact NMs in other domains, in case the subscriber's traffic traverses multiple domains.

‡NSIS stands for Next Steps in Signaling.

§PIB is a logical entity. It can be part of the Policy Decision Point (PDP).

If the NM is not able to allow the requested service for some reason (such as due to lack of resources or if the subscriber is not authorized to enjoy the requested service), a negative acknowledgment is sent to the subscriber. The negative acknowledgment may also include the reason for turning down the request, and a list of services that the subscriber can currently subscribe to. Including such information might induce the subscriber to negotiate again a different service, while otherwise the subscriber would have abandoned negotiations. Once the service has been established, a subscriber can re-negotiate for a different service at a later instant. A procedure similar to the above is followed for modifying an already existing service.

While the subscribers usually initiate the negotiation, it is also possible for the network to initiate the negotiation under some conditions. For example, if resources in the network become scarce, the NM can negotiate with the subscribers requesting them to degrade their existing SLSes to suit the current network conditions. The NM could offer cost incentives to the subscribers who accept the suggested SLS. Similarly, when there are unused resources available, the NM could also offer better services at a lower price to the subscribers.

The above discussions assume that the entities NM and RM are centralized in nature. Many of the negotiation protocols [3, 6, 7] work under this assumption, leading to scalability considerations. It is to be observed that this architecture separates signaling traffic from data traffic and the NM and RM are responsible for handling only the signaling traffic. When compared to data traffic, the signaling traffic is much smaller and hence a centralized architecture should be able to handle a large number of subscribers. Also, service negotiation does not take place each time the service is to be used, but only when the offered services or the subscriber needs change, creating the need for new agreements to optimize the new situation. Such centralized controllers have been successfully employed in other IETF protocols such as Megaco (IETF RFC 3015), COPS (IETF RFC 2748), and Middlebox (IETF RFC 3303). There are also negotiation protocols [5, 8] that work with distributed implementations of NM and RM. The working of such protocols is similar to the one discussed before but for the fact that, instead of a single NM and RM, the routers present along the traffic path collectively function as NM and RM.

Having discussed the general behavior of the negotiation protocols, we will now study the characteristics of each of them.

3.2 Resource Negotiation And Pricing (RNAP) protocol

RNAP [5] was one of the earliest protocols developed to facilitate dynamic service negotiation in NGI and can be considered as an extension of RSVP. A distinguishing characteristic of RNAP is that, apart from negotiating SLS parameters, it also allows for the negotiation of price of the contracted services. It can work with both centralized and distributed implementations of NM and RM. Another feature of RNAP is that it employs a *soft state* approach for negotiation and hence periodic signaling from the subscriber is required to *refresh* the negotiated services.

3.3 Service Negotiation Protocol (SrNP)

Service Negotiation Protocol (SrNP) [6] is a protocol developed by the Tequila [2] consortium. A unique feature of SrNP is that the protocol is not specific to any particular SLS format or to the context of a SLS. It is general enough to be applied for negotiating any document provided the document is in the form of attribute–value pairs. The semantics and format of the document under negotiation are transparent to the protocol. The objective of the negotiation process is to agree on the value of the attributes included in the document under negotiation, rather than the attributes themselves. SrNP also allows for the NM to put a hold on the request, i.e. to postpone its response to the subscriber's request, should the server's negotiation logic sees that an agreement is likely to be reached in the near future. SrNP messages could be encoded in ASCII, BER/TLVs or XML as convenient for the stack used. It is also possible to encapsulate SrNP messages in widely deployed protocols such as RSVP (by defining new TLVs) and COPS (by specifying a new client-type).

3.4 COPS-SLS

COPS-SLS [7] is an extension of COPS protocol [10] to negotiate a service level specification either between a customer and a network or between two networks. A characteristic feature of COPS is that it distinguishes the interactions between a subscriber and the NM into two phases: *configuration* and *negotiation*. In the configuration phase, the service provider informs the subscriber how to request a level of service. For example, it supplies the client with information about the negotiation mode and the time interval to renegotiate. After successfully installing the configuration, the subscriber can start the negotiation phase.

3.5 Dynamic Service Negotiation Protocol (DSNP)

Dynamic Service Negotiation Protocol (DSNP) [3] is a protocol developed by the ITSUMO (Internet Technologies Supporting Universal Mobile Operation) team from Telcordia Technologies, Inc. and Toshiba America Research, Inc. (TARI). Unlike RNAP, SrNP, and COPS-SLS that are extensions of other existing network protocols, DSNP was developed exclusively for dynamic service negotiation. Consequently, it is light-weight and is better suited for wireless devices such as PDAs and mobile phones. Also, by its architecture it provides better support for wireless mobile subscribers. Whenever a mobile subscriber negotiates for some service, the NM disseminates the QoS profile of the subscriber not only to the edge router that serves the wireless network in which the subscriber is currently located, but also to those that serve the wireless networks that are adjacent to the subscriber's current location. Consequently, when the subscriber moves into any of its adjacent networks, he/she continues to enjoy the negotiated service without any additional signaling.

3.6 QoS NSIS Signaling Layer Protocol (QoS-NSLP)

QoS-NSLP is a protocol for signaling QoS reservations in the Internet [8]. It is currently being defined by the IETF's NSIS working group and is an extension of RSVP[¶]. Its working is quite similar to that of RNAP and it too employs a soft-state approach for service negotiation. QoS-NSLP does not assume the existence of a centralized resource manager in each domain to carry out the negotiation process. In fact, the protocol is a better fit for distributed settings.

3.7 QoS Generic Signaling Layer Protocol (QoS-GSLP)

QoS-GSLP [9] is a protocol proposed by the Ambient networks consortium and is a part of the Generic Ambient Networking Signaling (GANS) protocol suite. It is used for controlling and negotiating bi-lateral QoS requirements in mobile wireless environments and builds on top of the IETF's NSIS protocol suite. It employs path de-coupled signaling and uses knowledge about mobility patterns and traffic patterns to

[¶]QoS-NSLP was primarily designed as a protocol for signaling QoS requirements. On account of the protocol's generic design, its ability to support service negotiation has been discussed in NSIS charters. Hence, QoS-NSLP has been included in the discussions.

set up SLSes well in advance. This reduces the SLS set-up times for mobile subscribers, which can be advantageous in a dynamic environment.

4 Characteristics of a service negotiation protocol

This section outline the features that are desired in an ideal service negotiation protocol. Based on these features, we will then draw a comparison between the protocols discussed in the previous section.

1. *Primary negotiation capabilities:* Any negotiation protocol should support the following set of operations that are vital for negotiating and establishing a SLS dynamically [2]:

   (a) A client should be able to specify and request a new service with its service provider.

   (b) A service provider should be able to communicate its acceptance or rejection of the requested service to the client.

   (c) The protocol should enable a service provider to modify a requested service and re-negotiate with the corresponding client.

   (d) A client should be able to accept or reject a service proposed by the service provider.

   (e) The service provider should be able to modify a client's accepted service if need be.

2. *Compatibility with QoS architectures:* In NGI, traffic between the end hosts may have to pass through networks owned by several service providers. Therefore the protocol employed for end-to-end service negotiation should be compatible with most, if not all, standard QoS architectures.

3. *Transparency to wireless link layer technologies:* In NGI, traffic between two hosts may have to pass through networks employing different wireless link layers. Hence it is required that any protocol aiming to support end-to-end service negotiation should be compatible with all layer-2 wireless technologies. Such a protocol will also provide the needed transparency for a subscriber moving between disparate wireless environments.

4. *Reduced Signaling Overhead:* The negotiation protocol should be scalable in terms of the signaling required between the subscriber and the service provider. For example, consider a mobile node enjoying a certain service and which has moved to an adjacent network/cell with enough resources. The

mobile node should not be required to negotiate for the same service again, just because it has moved to a new network/cell. Signaling consumes precious resource like bandwidth and battery power of a wireless device, and hence should be sparingly used.

5. *Transparency to SLS parameters:* All the service negotiation protocols are aimed at negotiating the parameters of a SLS. As discussed under section 2, though a consensus can be arrived at as to what set of parameters constitutes a SLS, the IETF is yet to standardize the format of a SLS. Consequently, it is advisable that any service negotiation protocol that has been/is being developed does not pre-define the format of a SLS.

6. *Extend existing protocols:* A service negotiation protocol, preferably, should make use of existing signaling protocols. This is because, already several routing protocols and network management protocols run on top of present day networks. Adding one more protocol for service negotiation may increase the complexity of network administration.

7. *Light-weight:* Mobile subscribers with ubiquitous connectivity will benefit to a great extent on account of dynamic service negotiation. Consequently, the service negotiation protocol is most likely to be used across devices with varying capabilities in terms of battery, computing power, and memory. Therefore, it is imperative that the protocol should be light-weight. Extending protocols that are originally developed for other purposes to do service negotiation might become too complex for a mobile device, when compared to a protocol dedicated just for this objective. Note that this requirement is in contrast to the one discussed under item 6.

In the above list, items (1) to (5) are design features while (6) and (7) are implementation characteristics.

5 Comparative study of the negotiation protocols

Almost all the protocols possess the first three characteristics presented under section 4. All of them support primary negotiation capabilities and follow widely accepted QoS architectures. They all carry out the negotiation at the IP layer and thus are transparent to the wireless link layer technologies as well. While all the protocols support the first three characteristics, they score differently with respect to the remaining four.

Apart from DSNP, all negotiation protocols are extensions of existing network protocols. RNAP and QoS-NSLP are extensions of RSVP, while COPS-SLS is an extension of COPS. Though QoS-GSLP is a part of the GANS protocol suite, it is expected to be backward compatible with QoS-NSLP. SrNP can be implemented by extending either RSVP or COPS. Thus, all protocols barring DSNP, may not add any administrative complexity/overhead to the network in terms of introducing a new protocol. While DSNP has this drawback, it wins over other protocols in terms of its light-weight and reduced signaling overhead. As a new protocol, DSNP was developed exclusively for service negotiation. This focused approach enables DSNP to be light-weight. QoS-GSLP also was developed exclusively for service negotiation and hence is light-weighted too. RSVP and COPS were originally developed to achieve different goals. While these protocols can be extended to provide service negotiation capabilities, the broad scope of the resulting negotiation protocols will increase the overall complexity. On account of their light-weight, DSNP and QoS-GSLP might be more suitable for devices with limited capabilities such as PDAs and other hand-helds than RNAP, SrNP, COPS-SLS, and QoS-NSLP.

In addition to the complexity, there are message overhead considerations: protocols such as RNAP and QoS-NSLP use soft-state which requires periodic signaling to refresh the negotiated service. In wireless devices, such periodic signaling consumes precious resources such as battery power and bandwidth. On account of this reason too, RNAP and QoS-NSLP may not be suitable for wireless networks. DSNP, COPS-SLS, and SrNP do not require periodic signaling. DSNP and COPS-SLS can also make use of their network architectures to reduce the signaling load for mobile subscribers when they move into new access networks. With QoS-GSLP, there is a possibility of additional communication overhead to maintain bi-lateral QoS agreements and extra complexity to handle end-to-end services.

In terms of SLS transparency, SrNP might be a better choice than others. This is because SrNP does not assume any specific format for the SLS. Other protocols adopt specific formats for the SLS that is being negotiated[‖] in spite of there being no universally accepted format for SLS specification yet. On account of this transparency, SrNP messages tend to be bigger than other protocols, thereby increasing the signaling overhead. The above discussions are summarized in table 1.

[‖] This comparison does not take QoS-NSLP into consideration, as the detailed specifications of QoS-NSLP for service negotiation are yet to be defined.

6 Conclusions

This paper presents an overview of the various service negotiation protocols that have been proposed so far for next generation Internet. Though all the protocols possess the basic negotiation capabilities, they differ in terms of their signaling costs, protocol complexity, SLS transparency, and administrative overhead. No single protocol emerges as a clear winner with respect to the aforesaid parameters. Currently, there are some on-going initiatives including Europe's EuQoS project and IETF's path-decoupled signaling. They potentially could be used for service negotiation too. While the evolution of the discussed protocols indicates the progress made in the field of dynamic service negotiation, more work needs to done. Specifically, additional research is required to address the following:

1. While the specifications of the protocols have been done, none of them have been widely deployed on a real network. Therefore, the issues involved in their deployment, if any, need to be looked into. A related concern would be to resolve the issues involved in inter-working of these negotiation protocols with other IETF protocols.

2. Dynamic service negotiation can impact other network functionalities such as QoS routing, and network provisioning. The inter-play between dynamic service negotiation and other network components needs to be further investigated.

3. Today, many different wireless systems exist, ranging from wireless Personal Area Networks (PANs), wireless Local Area Networks (LANs) to outdoor cellular systems. It is widely believed that this incompatibility will continue to exist in the future. An important question that arises is how to adapt these service negotiation protocols for a heterogeneous wireless setting such as 3G, 4G, and ad hoc networks.

4. So far, a service is negotiated mostly on the basis of network resource availability against the service QoS requirements. In the future, research will focus on developing more sophisticated mechanisms to reason on security issues (security levels, encryption algorithms, security keys, authentication protocols) and to dynamically formulate the service prices using econometric models. The negotiation protocols should support the requirements of these new mechanisms.

References

[1] D. Dobler and D. Burt, *Purchasing and Supply Management*. New York: McGraw Hill, 6 ed., 1996.

[2] D Goderis et al., “Service level specification semantics, parameters and negotiation requirements.” IETF Internet Draft, <draft-tequila-sls-01.txt>, work in progress, June 2001.

[3] J.-C. Chen, A. Mcauley, V. Sarangan, S. Baba, and Y. Ohba, “Dynamic Service Negotiation Protocol (DSNP) and Wireless Diffserv,” in *Proc. of IEEE International Conference on Communications (ICC)*, (New York, NY), pp. 1033–1038, April 2002.

[4] M. D’Arienzo, A. Pescape, and G. Ventre, “Dynamic service management in heterogeneous networks,” *Journal of Network and Systems Management*, vol. 12, September 2004.

[5] X. Wang and H. Schulzrinne, “RNAP: a resource negotiation and pricing protcol,” in *Proc. of International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV)*, (Basking Ridge, NJ), pp. 77–93, June 1999.

[6] Tequila Consortium, “SrNP: Service Negotiation Protocol.” http://www.ist-tequila.org/deliverables (last accessed on April 20, 2005), Oct. 2001.

[7] T. M. T. Nguyen, N. Boukhatem, Y. G. Doudane, and G. Pujolle, “COPS-SLS: a service level negotiation protocol for the Internet,” *IEEE Communications Magazine*, pp. 158–165, May 2002.

[8] S. V. den Bosch, G. Karagiannis, and A. McDonald, “NSLP for Quality-of-Service signalling.” IETF Internet Draft, <draft-ietf-nsis-qos-nslp-06.txt>, work in progress, Feb. 2005.

[9] Ambient Networks consortium, “Connecting Ambient Networks - Architecture and Protocol Design (release 1).” Deliverable D 3.2, Mar. 2005.

[10] D. Durham, Ed., J. Boyle, R. Cohen, S. Herzog, R. Rajan, and . Sastry, “The COPS (common open policy service) protocol.” IETF RFC 2748, Jan. 2000.

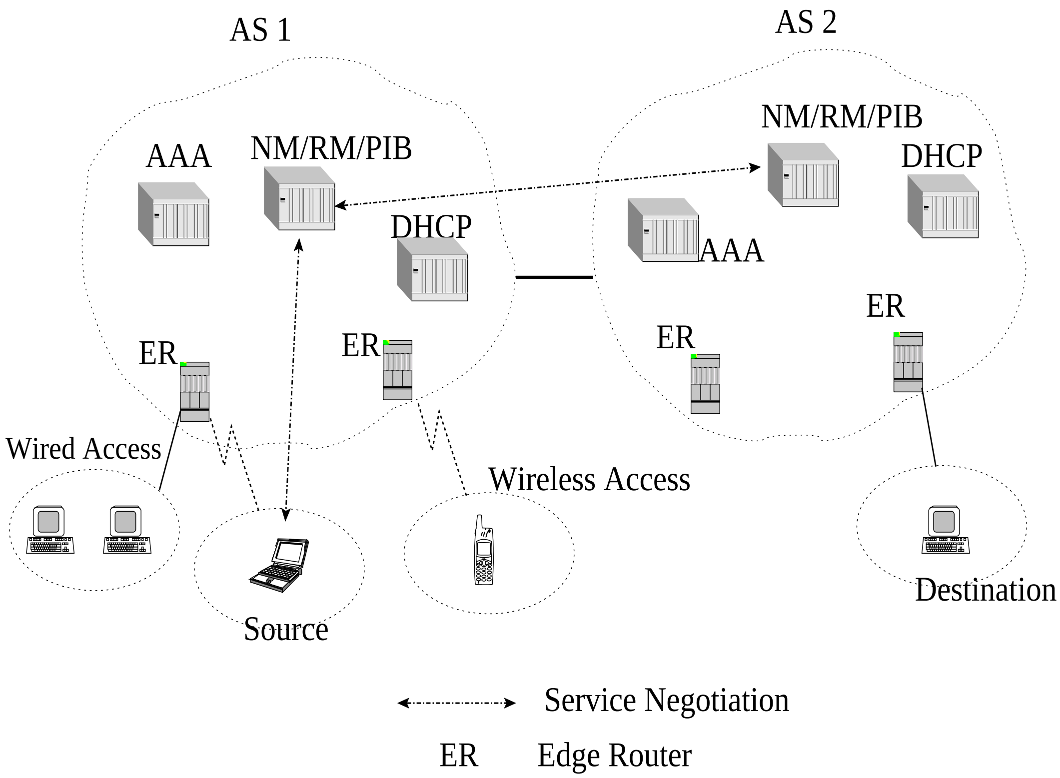

Figure 1: Network Architecture for service negotiation protocols.

	Protocol					
Feature	RNAP	SrNP	COPS-SLS	DSNP	QoS-NSLP	QoS-GSLP
Primary Negotiation	Yes	Yes	Yes	Yes	Yes	Yes
Generic QoS Architecture	Yes	Yes	Yes	Yes	Yes	Yes
Link layer transparency	Yes	Yes	Yes	Yes	Yes	Yes
Extend existing protocol	Yes	Yes	Yes	No	Yes	Yes
Light-weight	No	No	No	Yes	No	Yes
Reduced signaling	No	May be	Yes	Yes	No	May be
SLS format Transparency	No	Yes	No	No	May be	No

Table 1: Comparison of service negotiation protocols.