A Distributed Policy Approach in Support of Multimedia Session Establishment

Hamid Syed Mahmood, Louis Nicolas Hamer
Nortel Networks, 100-Constellation Crescent, Nepean, Ontario. K2G 6J8.
E-mail: {hmsyed, nhamer}@nortelnetworks.com.

Abstract: Policies may be required in various scenarios during a multimedia flow set up, for example, during session setup, policy decisions must be provided to ensure that the session being requested lie within the bounds of the service profile established for the requesting host. Similarly, when a host requests resources to provide a certain QoS for a packet flow, policies may be provided to ensure that the required resources lie within the bounds of the resource profile established for the requesting host. The policy framework at the IETF [1] is focused on the resource reservation issues and thus provided a centralized approach for the policy decision function. This paper attempts to highlight the various scenarios in a multimedia flow establishment that leàd to a distributed approach for policy decision-making. The paper also discusses the required characteristics of an interface between the policy decision functions residing in different administrative domains.

Key words: Distributed Policy, policy decision function, service policy, and resource policy

1. INTRODUCTION

Policies may be required in various scenarios during a multimedia flow set up, for example, during session setup, policy decisions must be provided to ensure that the media streams being requested lie within the bounds of the service profile established for the requesting host. Similarly, when a host requests resources to provide a certain QoS for a packet flow, policies may

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: 10.1007/978-0-387-35620-4_43
D. Gaïti et al. (eds.), *Network Control and Engineering for QoS, Security and Mobility*

be provided to ensure that the required resources lie within the bounds of the resource profile established for the requesting host. The resource allocation protocol (RAP) WG at the IETF defined a policy framework (RFC 2753) [1] to provide a policy-based admission control. The framework focuses mainly on admission control issues on usage of the network resources. The role of the policy decision point (PDP) is discussed for resource admission control only in the context of RSVP and QoS provisioning. Extending the policy-provisioning context for scenarios other than resource admission control opens issues that are not in RFC 2753. The framework described in [1] is more a centralized approach where a policy decision-making device acts as the Policy Decision Point (PDP) for controlling all kinds of requests generated within the network. Policies are centrally stored and can be accessed by the PDPs. Due to the variety of possible business models and environments (explained later in this paper), a single policy decision entity providing policy decisions for both service control and resource control is not a practical approach. The PDP functionality needs to be distributed based on the type of policy decisions required and the kind of the business model employed.

The paper is organized as follows: Section 2 introduces the terms used in this paper. Section 3 provides a quick view of the centralized policy framework and describes the various scenarios where the policy decisions needed to support a multimedia flow should be made by entities that could either be in different domains or have different functional capabilities. The distributed policy management approach is then introduced in section 4. Section 5 presents an application of the concept proposed by this paper in UMTS networks. The paper ends with the conclusion of the concepts discussed.

2. DEFINITION OF TERMS

Most of the terminology used is from [1]. This section defines the few new terms used in this paper.

Access Provider (AP): The access provider is a logical grouping of elements that provide connectivity along the packet forwarding paths to and from an End Host. The AP contains network elements and policy decision-making entities whose responsibilities include management of resources along the packet forwarding paths.

Service Provider (SP): The Service Provider is a logical grouping of elements that offer applications and content to subscribers of their services. The Session Management Server (also know as application proxy) resides in the SP.

Administrative Domain: An administrative domain is the part of the network that is controlled by a single set of policies. One SP or AP owns an administrative domain.
Service Policies: The policies enforced by the service provider. These policies describe the user privileges in accessing their services.
Resource Policies: These policies describe the rules for accessing the network resources such as the resources allocated for a particular type of service on an interface, the mapping of subscriber flow requirements to the actual resource requirements etc.
Policy Decision Function: The policy Decision Function is a logical entity whose responsibility is to process the requests against the domain's policies and provide policy decisions.

3. THE POLICY DECISION-MAKING FOR A MULTIMEDIA SESSION ESTABLISHMENT

In most networks, each and every request to the network is compared against the network policies. The network operators set their objectives in the form of policy rules that are enforced and provided to every request that is generated within the network. Establishing multimedia sessions must take into account requirements for end-to-end QoS, authorization of network resource usage and accurate accounting for resources used. Network policies define the necessary rules to ensure that the multimedia sessions are within the permissible bounds. The policy decisions for a multimedia session does not come in a single step but are provided by the network in various phases of the session establishment. To set the context for the whole discussion of centralized vs. distributed policy system, lets take a brief look at the various kind of policies and policy decisions that are required for a multimedia session setup with QoS. The policy decision-making starts right from the instant when an end host attempts to access the IP network. The first policy decision comes from the authentication server that authenticates the end user for accessing the network. The decisions normally come from the AAA server and the network also provides the host an IP address (in case a DHCP server is involved). Such policies can easily be termed as "IP access policies". Along the path of a communication, there may be administrative entities that need to impose policy constraints on entities such as security gateways and router filters. There also is a need for end-points of a security association and/or, for their respective administrative entities, to securely discover and negotiate access control information for the end hosts and for the policy enforcement points (security gateways, routers, etc.) along the path of the communication. The "IP security policies" (also an IETF

Workgroup: IPSP) provides the means for managing IPSEC security policy, negotiating security association (SA) parameters between IPSEC endpoints, and distributing authorization and policy information among hosts that require the ability to communicate via IPSEC. During a multimedia session establishment, the end user applications negotiate the session description parameters for an agreed level of QoS end-to-end. Lets assume that the end user attempts a VoIP session using the session initiation protocol (SIP [8]). The SIP request is used to negotiate the session description parameters (particularly the codec information). A SIP server (Session Management Server) in the network first authorizes the request parameters against the network policies stored in the users profile. A decision comes from the network whether or not the request, to establish a SIP session for a VoIP call, can be permitted. Such decisions are termed "service policies".

Finally, the access to the network resources and their usage are controlled through the "resource allocation policies". The resource allocation policies enable the network operator to manage and control the access to their network resources in a scalable and reliable fashion. The policy information structures to manage a Differentiated Services network and integrated services [9] using the Resource Reservation Protocol are a good example of resource allocation policies.

3.1 Why Centralized Policy Management Systems are not sufficient

The diagram in Figure 1 provides a simplistic view of the elements that encompass a policy-enabled network. All the elements of the policy management system are within the same administrative domain. A policy server is responsible for providing policy decisions to PEPs by referencing information stored in a Directory Service. Since the centralized policy management approach assumes a single administrative domain to enforce network policies, the policy server is responsible for providing the policies to the PEP. These policies include the rules for the IP access, Security, Session authorization and resource allocation. The centralized policy server approach only works in a scenario where the providers of the services and the resource access lie in the same administrative domain. However, there could be a number of scenarios where the centralized approach may not work. Moreover, even within the same administrative domain, it could be more efficient and beneficial to distribute the policy decision function to separate policy servers. In the following sub sections, we discuss some scenarios where the policy decision-making function may be distributed to separate policy servers. This applies to a single administrative domain as

well as to the multiple domains controlled by different service providers and access providers.

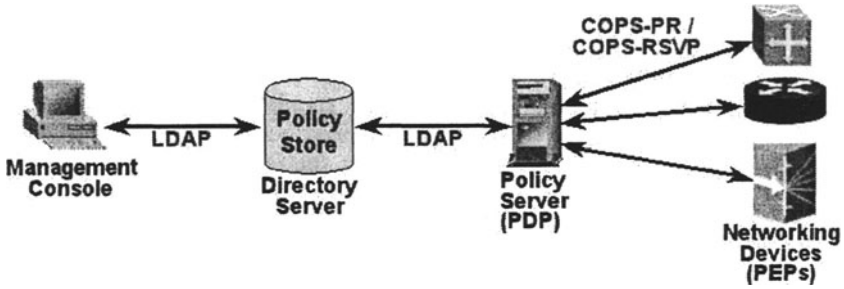

Figure 1. Centralized Policy Management Model

3.1.1 Distributed Policy Decision Function within same Administrative Domain

A single service or access provider may distribute the policy decision-making function based on the type of the policy decision required to separate policy servers. For example, the policy decisions for security, resource allocation and authorization may be outsourced to separate assigned policy servers in the same domain. The policy servers may share a single database or each may employ its own policy database. In either case, the policy servers may need to interact with each other to verify any previously made decision. For example, the policy server responsible for the session authorization policies may have granted a policy decision for the user authorization for a multimedia session request. During the resource allocation phase, the policy server for resource allocation will need to interact with the session authorization policy decision function for confirmation and to avoid potential fraud.

3.1.2 Separate Providers for Services and Resources

This scenario covers the situation where the service provider and the resource provider are separate business entities. They control different policies, and their interactions are required for end-end multimedia session management. Consider there is a policy server (PS) in the Resource Provider domain that is separate from the PS in the Service Provider domain. Policies are distributed between both policy servers – the service policies reside in the Service Provider PS and the resource policies reside in the Resource Provider PS. A Framework for session setup with media authorization [2],

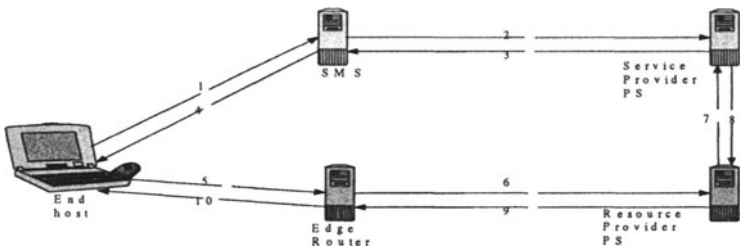

Figure 2. Separate Providers for Services and Resources Scenario

describes different models, two of which, the associated model with two policy servers and the non-associated model, present detailed walkthroughs where service policies and resource policies are distributed. For better readability of this paper, a walkthrough, of the associated model with two policy servers, is presented below. Figure 2 illustrates the scenario.

1.The End Host issues a session set-up request (e.g. SIP INVITE) to the Session Management Server indicating, among other things, the media streams to be used in the session. As part of this step, the End Host may authenticate itself to the Session Management Server.

2.The Session Management Server, possibly after waiting for negotiation of the media streams to be completed, sends a policy decision request (e.g. COPS REQ) to the SP Policy Server in order to determine if the session set-up request should be allowed to proceed.

3.The Policy Server in the session control domain applies his "service policies" by referencing the users service profile. The Policy Server sends a decision (e.g. COPS DEC) to the Session Management Server, possibly after modifying the parameters of the media to be used. Included in this response is an "authorization token" that can subsequently be used by the SP Policy Server to identify the session and the media it has authorized. The authorization token contains a field identifying its generator.

4.The Session Management Server sends a response to the End Host (e.g. SIP 200 or 183) indicating that session set-up is complete or is progressing. Included in this response is a description of the negotiated media along with the authorization token from the SP Policy Server.

5.The End Host issues a request (e.g. RSVP PATH) to reserve the resources necessary to provide the required QoS for the media stream. Included in this request is the authorization token from the SP Policy Server provided via the Session Management Server.

6.The Edge Router intercepts the reservation request and sends a policy decision request (e.g. COPS REQ) to the RP Policy Server in order to determine if the resource reservation request should be allowed to proceed.

Included in this request is the authorization token from the SP Policy Server provided by the End Host.

7.The RP Policy Server uses this token to learn which SP Policy Server authorized the media. It then sends an authorization request to that SP Policy Server in order to determine if the resource reservation request should be allowed to proceed. Included in this request is the authorization token from the SP Policy Server provided by the End Host.

8.The SP Policy Server uses this token to correlate the request for resources with the media authorization previously provided to the Session Management Server. The SP Policy Server sends a decision to the RP Policy Server on whether the requested resources are within the bounds authorized by the SP Policy Server. Alternatively, the correlation could be done at the RP Policy Server.

9.The RP Policy Server applies his "resource policies" and sends a decision (e.g. COPS DEC) to the Edge Router, possibly after modifying the parameters of the resources to be reserved.

10. The Edge Router, possibly after waiting for end-to-end negotiation for resources to be completed, sends a response to the End Host (e.g. RSVP RESV) indicating that resource reservation is complete or is progressing.

4. DISTRIBUTED POLICY MANAGEMENT

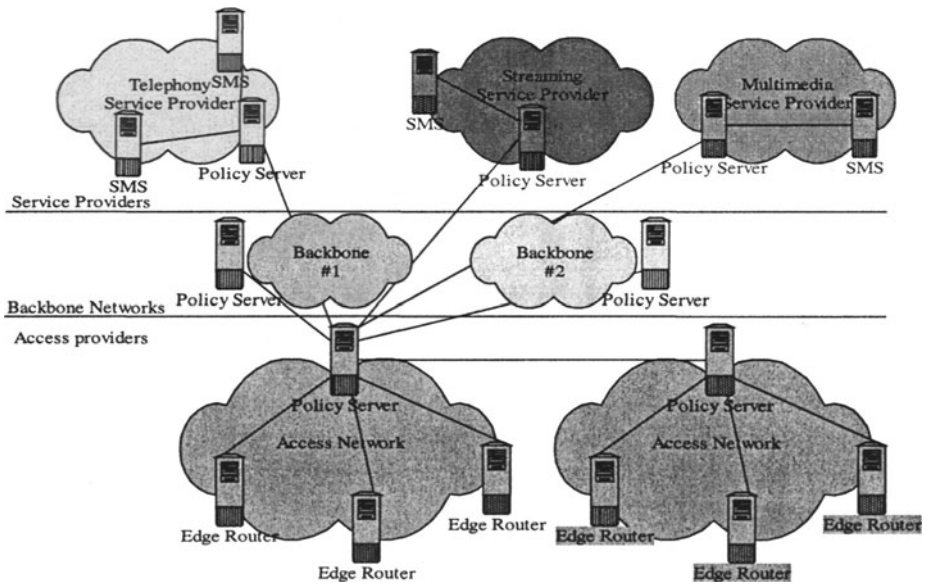

*Figure 3.*Distributed Policy Model

Figure 3 illustrates the different scenarios introduced in section 4 and gives a realistic picture of today's Internet. The Internet is a collection of operators providing services. These services include Internet access and connectivity, backbone transport, telephony, multimedia, etc. The Internet is a group of operators specializing in one or more of these services. A successful operator will forge agreements with others to provide their customers a variety of services. The policies are most likely distributed across different domains operated by single or different operators. The centralized or the single policy server approach may not work for these scenarios where the multimedia call has to request policy decisions from more than one administrative domain. In fact, it may open issues of fraud and inaccurate billing. The policy management for such networks is a distributed model with policies that are distributed across various domains. In this section, we will discuss the various elements (and interfaces between the elements) of the network required for the distributed policy management approach. We aim to highlight the differences from the centralized policy management and identify the open issues that may require further research and standardization. The various elements of the distributed policy framework (shown in Figure 3) are as follows:

- The policy clients or policy enforcement points (PEPs) are functionally located at the edge of the access network. These edge routers needs to enforce the policies discussed in section 4 to allow/establish a multimedia call. The client would download the policies that can be provisioned by the policy servers located in the same admin domain or in a different domain depending upon the nature of the policy required.
- The policy servers are distributed across different domains and are responsible for the policy provisioning for specific functions or services supported by their respective domains.

The distributed policy framework does not impact the policy data modelling being taken at the IETF policy framework WG [10]. In fact the various data models for QoS, security and accounting are focused on specific policy functions independent of what policy management model uses these policies. For example, these functions may all be supported in a single domain (centralized policy approach) or may be distributed across various policy servers located in different domains (distributed policy approach). The policy client (PEP) and the policy server (PDP) within the same domain may use the standard common open policy services (COPS) protocol (defined by the IETF) [6] for the policy transactions. The policy decision function interactions are the key to the distributed policy management. Such interactions are important to support the scenarios where a user wishes to setup a multimedia session spawning across different providers or policy domains. The characteristics of an interface for policy decision function

interactions have not been discussed in detail yet although some interesting research has started in the AAAArch WG [3,4] at the IRTF. The Common Open Policy Services (COPS) has also been proposed as a way to perform the service level negotiation between the customer and service provider networks [11]. While COPS appears to be a suitable protocol to be proposed for policy server interactions (mainly because it is meant to carry policy information) however there are a few characteristics of COPS that may not be suited particularly for policy exchange between the boundaries of two different administrative domains. The connection-oriented nature of the COPS transport and the keep-alive messages seems to be overhead for policy server interaction across domains. On the other hand, reliability is key in policy exchange that a connectionless approach cannot provide. The issue of security needs to be well understood for such an interface that requires policy exchange between domains owned by different operators.

To further explain the policy server interaction and the kind of policy exchanges required, a simple example is presented that is based on the session authorization model described in [2]. It will present to the reader a high level view of what kind of information may be required when the policy servers in access provider and the service provider networks interact for a SIP initiated VoIP session with the correct QoS required by the session.

4.1 An Example

Service providers offer multimedia to their customers. Most of these services demand a particular Quality of Service, e.g. low delay and latency. In order to meet the users QoS demand, the service provider must be able to request the access provider to provide this particular users session with QoS.

4.1.1 Authorization Request

The access policy server just received a COPS request. In this request is an authorization token identifying the service policy server who authorized the service and a session identifier. The access policy server will attempt to contact the service policy server to request authorization for the session. This message must contain the session identifier. Upon receipt of this message, the service policy server will use the session identifier to retrieve the details of the authorized session. A subset of those details could be for example: VoIP Session, codec g.711, from IP address W.X.Y.Z to IP address A.B.C.D.

4.1.2 Authorization Decision

The details of the session stored in the service policy server could then be easily formatted into an RSVP flowspec, for example, and returned to the resource policy server. The next step would be to compare both the authorized flowspec received from the service policy server and the flowspec received in the resource request of the end host, to make sure the requested QoS is within the bounds of the authorized QoS.

5. APPLICATION OF THE CONCEPT TO UMTS NETWORKS

The concept of distributed policy is present in UMTS networks in the 3GPP Release 5. As shown in figure 4, policies are distributed between the IM Subsystem and the UMTS Core Network – please note not all network entities are shown for simplicity. Consult [14], [15] & [16] for more details. For readability purposes, some 3GPP term definitions are provided:

IM Subsystem: The IP multimedia subsystem enables operators to offer their subscribers' multimedia services based on and built upon Internet applications, services and protocols.

P-CSCF: Proxy-Call Session Control Function - The P-CSCF is a network element providing session management services (e.g. telephony call control).). The P-CSCF is characterised by being the first contact point for the User Equipment within the IM subsystem (IMS)

S-CSCF: Serving-Call Session Control Function - The S-CSCF is a network element providing session management services (e.g. telephony call control). The S-CSCF actually handles the session states in the network.

HSS: Server containing the subscription-related information to support the network entities actually handling calls/sessions.

PCF: Policy Control Function – A PDP as defined in [6].

GGSN: Gateway GPRS Support Node - The GGSN is a network element connecting the User Equipment to the external network. The GGSN contains a PEP to enforce policies.

The Release 5 3GPP policy Architecture is based on the service based policies (e.g. the subscriber profiles) being stored in the HSS, which is part of the IMS. SIP requests are authorized at the HSS and relayed to the User equipment through the Proxy-CSCF. The P-CSCF, also plays a second role, which is to act as a proxy to the HSS to communicate the authorized session parameters to the PCF. The PCF, which stores resource based policies, can then correlate the authorized session parameters with the requested resources through the interface named Go, i.e. interface between the GGSN and PCF.

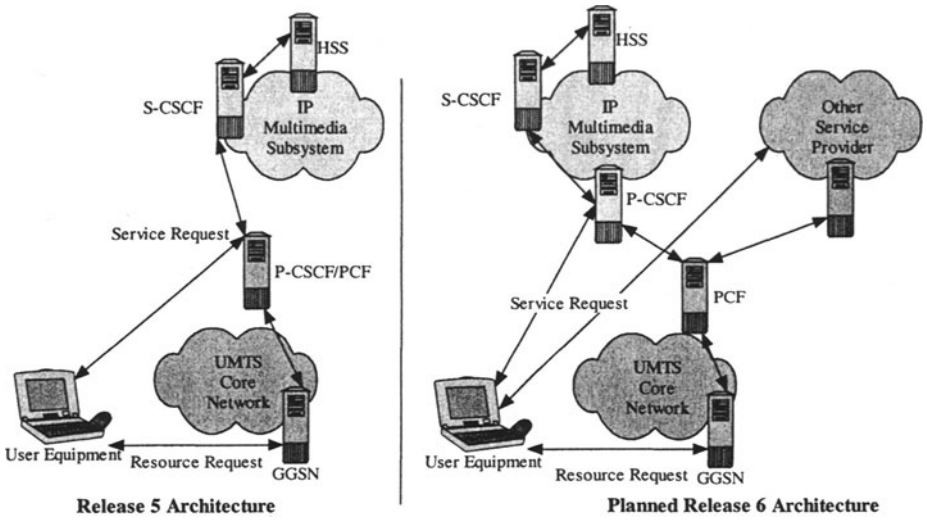

Figure 4: UMTS/3GPP Policy Architecture

Although the P-CSCF and PCF are two separate logical entities belonging to distinct domains (the IMS and Core Network domains), the release 5 3GPP specifications present the PCF as being an integral part of the P-CSCF. The main reason, time-constraint, led the participants of the 3GPP standardization body to couple both entities so that a standard interface did not have to be defined between them. The consequence of this close coupling is the limitation to the IP Multimedia Subsystem from requesting resources for their customers. In fact, the current architecture prevents other service providers from requesting resources. In release 6, a 3GPP work item has started to define a standard interface between the PCF and P-CSCF. This interface would allow the IM Subsystem from communicating authorization policies with the UMTS Core Network through the Proxy-CSCF. This interface would also enable any service provider to communicate policies with the UMTS Core Network. Since this is an application of distributed policies, a generic protocol should be defined allowing any service domain to provide session authorization to their customers in a particular access network. This open architecture would enable third party service providers to offer a multitude of services and even enable future killer-applications to access network customers. Past experience has proven this type of open architecture has been much more successful than garden-walled models such as the WAP system.

6. CONCLUSION

The Internet is not a homogenous network. On the contrary, it is composed of a multitude of administrative domains linked together. The current trend has been the logical grouping of network elements per administrative domain and per their specialization: access provider, backbone providers and service providers. Logical grouping has many benefits, e.g. separate evolution of each grouping and better manageability.

This structure brings, as described in this paper, situations where multiple administrative domains are involved in the setup of one multimedia session. For a successful setup of this session, cooperation between administrative domains is key. In this paper, we explained the scenarios where policy management require a distributed approach. We also discuss a comparison of distributed vs. a centralized approach in terms of functional elements and required interfaces. Since policies may be scattered in different policy servers, the need for a standard mechanism for policy servers to exchange and communicate their policies to others appears to be a major requirement in distributed policy management approach. Finally, although the concept of distributed policy management is applied in UMTS Networks, it was argued that a more open architecture would hopefully enable the seamless offering of new revenue-generating services.

7. REFERENCES

[1] R.Yavatkar et al., "A framework for Policy-based Admission control", IETF RFC 2753, January 2000.

[2]L-N. Hamer, B. Gage, "Framework for session setup with media authorization", Internet-Draft, draft-ietf-rap-session-auth-03.txt, June 2002, Work in progress.

[3] C. de Laat et al, "Generic AAA Architecture", IETF RFC 2903, August 2000.

[4] J. Vollbrecht et al., "AAA Authorization Framework", RFC 2904, August 2000.

[5] Y. Sunir et.al., "Policy QoS Information Model", draft-ietf-policy-qos-info-model-04.txt, November 2001, Work In Progress.

[6] D. Durham et al, " The COPS Protocol", IETF RFC 2748, January 2000.

[7] K. Chan et al., "COPS Usage for Policy Provisioning", RFC 3084, March 2001.

[8] Rosenberg et al., RFC 3261, "SIP: Session Initiation Protocol", June 2002.

[9] Y. Bernet et al, "A Framework for Integrated Services Operation over Diffserv Networks", IETF RFC 2998, November 2000.

[10] Policy Framework WG. http://www.ietf.org/html.charters/policy-charter.html

[11] T.M.T. Nguyen et al., "COPS Usage for SLS negotiation (COPS-SLS)", IETF Internet-Draft, draft-nguyen-rap-cops-sls-03.txt, July 2002, Work in progress.

[14] 3GPP TS 23.207: "End-to-End QoS Concept and Architecture", Release 5

[15] 3GPP TS 29.207: "Policy control over Go interface",Release 5

[16] 3GPP TS 23.228: "IP Multimedia (IM) Subsystem - Stage 2", Release 5