

Inter-Domain QoS Provisioning and Accounting

Tomoichi EBATA <ebata@sdl.hitachi.co.jp>

Masatoshi TAKIHIRO <takihiro@sdl.hitachi.co.jp>

Shigeru MIYAKE <yake@sdl.hitachi.co.jp>

Minoru KOIZUMI <m-koizu@sdl.hitachi.co.jp>

Hitachi, Ltd.

Japan

Felix HARTANTO <hartanto@fokus.gmd.de>

Georg CARLE <carle@fokus.gmd.de>

GMD FOKUS

Germany

Abstract

QoS[QOS]-capable communication network devices are now available, such as IP routers and switches, thus enabling the implementation of QoS-guaranteed services. Policy-Based Network Management (PBNM) is a framework to control these network devices with management rules that are called "Policy". At present, most network engineers on PBNM cope with only intradomain QoS-guaranteed services, but an enhanced framework for interdomain QoS-guaranteed services is required for real commercial networks that communicate with networks managed with a different policy. On real commercial networks, service providers have to charge customers for the interdomain QoS-guaranteed services.

We focus on interdomain QoS provisioning and accounting. First, we discuss the policy framework and its architecture. Based on this framework and architecture, we propose two protocols: a Policy Advertisement Protocol (PAP) that distributes local policies among domains, and a Policy Negotiation and Notification Protocol (PNP) that sets up a QoS-guaranteed communication path among domains[INTD].

Our experimental network confirms that this interdomain QoS provisioning is effective when service providers provide interdomain QoS services from the point of view of network administrator's tasks, performance, and network traffic.

Contents

- 1. Introduction
- 2. Interdomain policy framework
 - 2.1 Interdomain services
 - 2.2 Definition of policy
 - 2.3 Interdomain policy framework

- 2.4 Exchange policy
- 3. Interdomain policy exchange architecture
 - 3.1 Two types of architecture
 - 3.2 Subject
 - 3.3 Outline of policy exchange protocol
 - 3.4 Preparation
- 4. Interdomain PAP
 - 4.1 Outline
 - 4.2 Protocol
 - 4.3 Message format
- 5. Interdomain PNP
 - 5.1 Outline
 - 5.2 Protocol
 - 5.3 Message format
- 6. Evaluation
 - 6.1 Evaluation of PAP
 - 6.2 Evaluation of PNP
- 7. Conclusion
- 8. References

1. Introduction

QoS-guaranteed services can now be provided because of the development of communication network devices, for example QoS-capable routers. The Internet Engineering Task Force (IETF) has started to examine "policy" (sets of rules) as a means of managing and controlling network devices and providing services.

In addition, the use of Service Level Agreements (SLAs)[SLA] has been introduced for the benefit of Internet Service Providers (ISPs) and their users. These SLAs are contracts for services that ISPs provide their users. They contain information on QoS-guaranteed services and the accounting cost of the services. Both QoS-guaranteed services and the accounting are critical to the SLA as far as Internet services are concerned. Policy Servers (PSs) [TERM] that uniformly manage the above QoS-guaranteed services, accounting, and SLAs are being developed now. PBNM is the management method in which PSs control the network using policy.

Interdomain QoS-guaranteed services have become important, leading to a need for contracts that refer to different domains, e.g., different ISP networks, enterprise networks, and carrier networks. In the case of interdomain environments, because each domain is managed in accordance with a specific policy, it is natural that a policy in one domain is different from a policy in another domain. Therefore, it is impossible to seamlessly provide interdomain QoS-guaranteed services and accounting within the current PBNM framework. The Policy Working Group (WG) in the IETF has started to examine policy core schemas and QoS schemas that are currently being used in PSs. But the schemas are supposed to be used within intradomain networks.

The purpose of this paper is to show how to implement interdomain QoS-guaranteed services and interdomain accounting. To implement interdomain QoS-guaranteed services and interdomain accounting, network service providers, for example ISPs, must provide users with an interdomain QoS provisioning service and an interdomain accounting service as fundamental services.

At the very least, it must be possible to exchange policy among domains because one domain must know the policy of another domain. This paper focuses on the interdomain policy framework needed to implement this policy exchange process, the specific contents of policy and their properties, and network architectures suitable for policy exchanges. It also describes two protocols that can be used for this implementation: PAP and PNP.

This paper is organized as follows:

- Section 2 presents the interdomain policy framework and its application to interdomain QoS provisioning and interdomain accounting. This section also presents 10 different kinds of policy. Six of these are identified as relevant for interdomain QoS provisioning and accounting processes. At the end of this section, we explain why and how these policies are exchanged.
- Section 3 presents two architectures for policy exchange. One architecture uses a so-called super PS, while the other does not. The utility of these architectures is compared, and some problems with them are discussed.
- In Sections 4 and 5, the message formats and further details on the two protocols are explained.
- Section 6 presents the results and evaluation of an experimental network that has PSs that use the two protocols.

2. Interdomain policy framework

2.1 Interdomain services

In this paper, the interdomain services are composed of three kinds of services. Figure 1 shows building blocks of interdomain services.

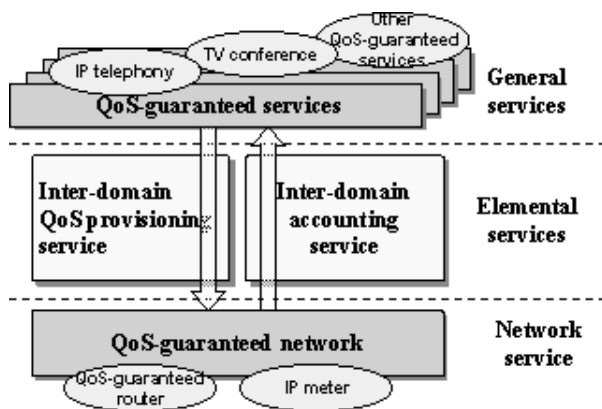


Fig. 2-1 Building blocks of QoS-guaranteed services

- General services: Services that customers use among domains, e.g., Internet Protocol (IP) telephony, TV conference. This paper does not address these services.
- Elemental services: Services that provide QoS provisioning and accounting for the general services among domains. These services are composed of policies that are defined in section 2.2
- Network services: Services that provide QoS-guaranteed networks. This services are composed of QoS-capable communication network devices, e.g., QoS-guaranteed routers and IP meters. This paper does not address these services.

2.2 Definition of policy

For the purposes of this paper, policies are defined as having attributes and rules that are composed of the attributes. They are similar to objects in the object-oriented technique (see figure 2).

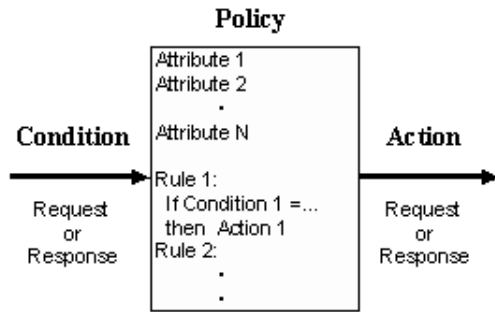


Fig. 2-2 Definition of policy

When events as conditions are input to the policy, the policy evaluates the events in accordance with the rules, and outputs the events as actions. The events contain request or response information.

2.3 Interdomain policy framework

The interdomain policy framework is shown in the figure 3.

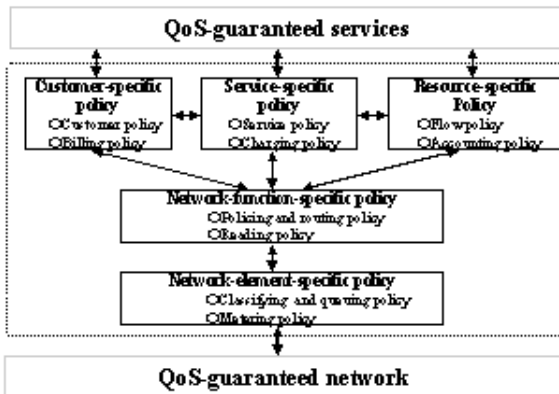


Fig. 2-3 Inter-domain policy framework

Customer-specific policy, service-specific policy, and resource-specific policy are to be exchanged with other domains. The main focus of this paper is the interdomain exchange of these policies.

Network function-specific policy and network element-specific policy are used to control the network devices for intradomain QoS provisioning and accounting.

The customer-specific policy controls the customer access to the available services within a provider's domain. The policy is composed of the attributes of the customer (e.g., priorities, usable services, valid

time, usable resources) and rules (e.g., "If Mr. A has gold priority, then he can use all services."). There are two kinds of customer-specific policy: customer policy and billing policy.

The service-specific policy controls how a service should be provided to a customer. The policy is composed of the attributes of the service (e.g., expected QoS, valid time, cost) and rules (e.g., "If service A is available, then its cost is one dollar a minute."). There are two kinds of service-specific policy: service policy and charging policy.

The resource-specific policy controls how actual network resources are made available. This policy is composed of the attributes of the resources (e.g., network medium, bandwidth, delay, jitter, Maximum Transmission Unit (MTU)), status information (kinds of service, path, and time), and rules (e.g., "If there is not enough bandwidth for Mr. A's service, reduce other customers' bandwidth"). There are two kinds of resource-specific policy: flow policy and accounting policy.

The network function-specific policy makes final decisions for each network device, or collects network status information from them. This policy is composed of the attributes of network functions and rules (e.g., "If the utilization rate of link A is above 90 percent, use link B"). There are two kinds of network function-specific policy: policing & routing policy and reading policy.

The network-elements policy defines the configuration of network elements for providing a network function associated with a service. The policy is composed of the attributes of network devices (e.g., setting parameters and their values) and rules (e.g., "If bandwidth is available, then set bandwidth parameter enable, and mark packets as green"). The policy is also composed of attributes of flow (e.g., real-time measurement values of bandwidth, delay) and rules (e.g., "If "user - IP address" is assigned, relate flow with IP address"). There are two kinds of network-elements policy: classifying & queuing policy and metering policy.

2.3.1 Interdomain QoS provisioning policy framework

Figure 4 shows the interdomain QoS provisioning policy framework, which is applied to the interdomain policy framework shown in figure 3. The QoS provisioning process starts with the requests from QoS-guaranteed services. This process is executed left-right and/or top-down.

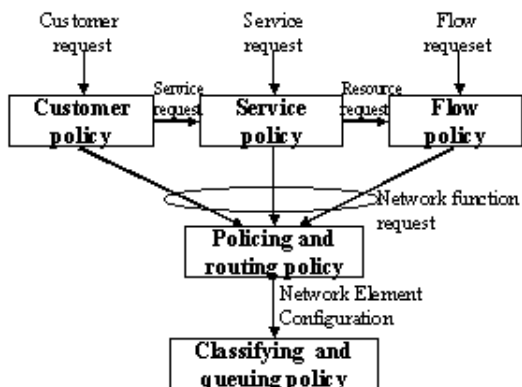


Fig. 2-4 Inter-domain QoS provisioning policy framework

The QoS-guaranteed services output the following information: (1) which service is needed,(2) who uses the service, and (3) how to use the service, in each of the customer policy, service policy, and flow

policy.

For example, when a QoS-guaranteed service puts out some requests, the customer policy checks the attributes of a customer and rules, and outputs the service request to the service policy. After checking the service attributes and rules, the service policy outputs the flow requests to the flow policy. Using the same process, the flow policy outputs the network function requests and the policing & routing policy dictates the classifying & queuing policy to be used to configure the communication devices.

2.3.2 Interdomain accounting policy framework

Figure 5 shows the interdomain accounting policy framework that is applied to the interdomain policy framework shown in figure 3. The accounting process starts with measurement for a specific network packet. This process is executed bottom-up and/or right-left .

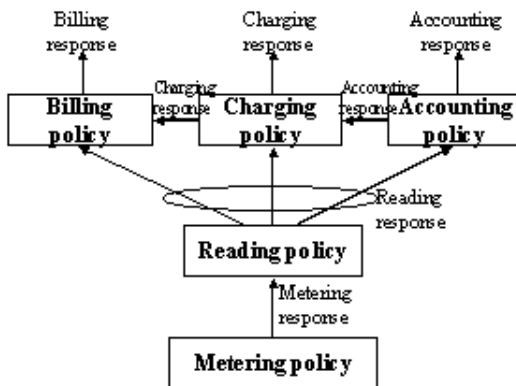


Fig. 2-5 Inter-domain accounting policy framework

For example, each communication device keeps observing network packets. The metering policy in a communication device checks the packet attributes and outputs the metering response to the reading policy. After checking the metering response, the reading policy outputs the reading response to the accounting policy and/or the charging policy and/or the billing policy. Using the same processes, the accounting policy outputs the accounting response and the charging policy outputs the charging response. Finally, the billing policy outputs the billing response in QoS-guaranteed services.

Thus, these responses are output to provide information for the following: (1) the customer's bill, (2) the status of the QoS-guaranteed service, and (3) the status of resource utilization in the QoS-guaranteed services.

2.4 Exchange policy

To execute interdomain QoS provisioning and accounting, a policy exchange is needed for the following reasons.

1. To provide interdomain services, it is necessary for one domain to know
 - (a) which services are available in another domain.
 - (b) which customers exist in another domain.
 - (c) how many resources are available in another domain.

2. To provide interdomain accounting services, it is necessary to have information on (a), (b), and (c), and
 - (d) additional information such as charging tables for services and resources in another domain.

It is also necessary to notify customers of the costs for used services and resources.

Thus, the policies of each domain involved in end-to-end service provisioning make it possible to provide interdomain QoS provisioning and accounting.

The above policies are exchanged in different ways because of the different property of each policy. Three ways of exchanging policies are described below.

2.4.1 Advertised policy

All relevant domains must have this policy so it can be carried by broadcast communication.

1. Interdomain QoS provisioning policy framework
 - (a) Customer policy
 - (b) Service policy
 - (c) Flow policy
2. Interdomain accounting policy framework
 - (a) Charging policy

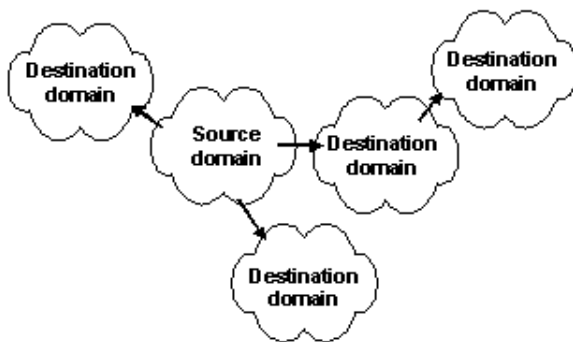


Fig. 2-6 Advertised policy

2.4.2 Notified policy

(1) Request and response

This policy is used as a trigger to start interdomain QoS provisioning. This policy may be exchanged by request and response communication, as shown in figure 6(a). A domain that sends requests is called a source domain and a domain that sends responses is called a destination domain.

Interdomain QoS provisioning policy framework

(a) Customer policy

(b) Service policy

(c) Flow policy

(2) Notification

This policy is used as accounting state reports to notify one domain of the cost of interdomain QoS-guaranteed service. This policy may be notified by unicast communication, as shown in figure 6(b). A domain that sends notifications is called a source domain. A domain that receives them is called a destination domain.

Interdomain accounting policy framework

(a) Billing policy

(b) Charging policy

(c) Accounting policy

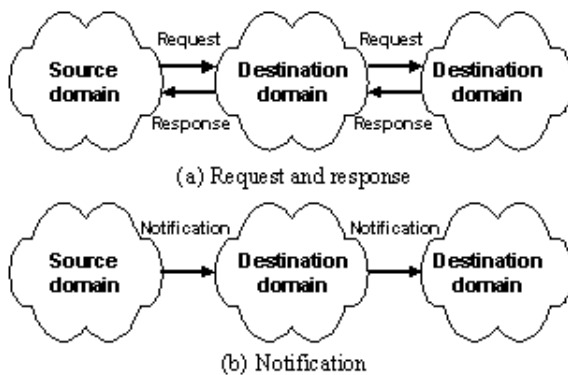



Fig. 2-7 Notified policy

3. Interdomain policy exchange architecture

In this section, the architecture of the interdomain policy exchange is explained.

3.1 Two types of architecture

Figure 7 shows two architectures.



Comparison

	Scalability	Flexibility	Machine Load	Policy Acquisition	Domain Administrator
(a)	Low	Low	Low(*1)	No Need	Need
(b)	High	High	High	Need	No Need

(*1) But SPS's load becomes high.

Fig. 3-1 Two types of architecture

In architecture (a), policies are exchanged through a Super Policy Server (SPS). A direct policy exchange process among domains is unnecessary here and the load on the PS becomes low. Because all PSs that execute QoS provisioning and accounting connect with the SPS directly, scalability and flexibility are sacrificed. And no one knows who manages the SPS.

In architecture (b), on the other hand, the load on the PS becomes high, but scalability and flexibility problems are not as serious as in architecture (a). Network administrators manage their intradomain only with the current Internet architecture.

3.2 Subject

The subjects of architecture of type (b) are as follows:

1. *Different exchange means for different properties of policy.* In section 2, it became clear that three communications, advertisement, request and response, and notification, are needed. So PSs need to have such means of communication.
2. *Low-load policy exchange process.* Because the six kinds of policy must be shared among domains, the load on PSs may be high. If the exchange processes are executed simultaneously, the processing power of PS may decrease. Therefore, PS should not exchange both the advertised policy and the notified policy simultaneously.

3.3 Outline of policy exchange protocol

To resolve the subjects described above, PSs should implement two phases of policy exchange processes.

1. *Policy advertisement phase.* In this phase, advertisement is executed to advertise policies among domains.
2. *Policy request-response and/or notification phase.* In this phase, request and response as well as notification are executed. After PSs refer to the advertised policy, PSs send and/or receive the policy that contains request and/or response contents. PSs send and/or receive the policies that contain the notification about the services too.

The PAP is used to advertise the policy. To support this functionality, BGP4 [RFC1364][BGP4] is

available. BGP4 is able to carry not only routing information but also additional information in a BGP4 message. In this paper, the enhanced BGP4 is proposed to advertise the policy of each domain.

The PNP is used to negotiate and notify domains of the policy. To support this functionality, COPS (Common Open Policy Service)[COPS], which is currently being standardized by the IETF, is available. COPS is a signaling protocol, and it exchanges policies between a policy decision point (e.g., policy server) and the policy enforcement point (e.g., router). This protocol is able to convey some policy and signal to execute some actions corresponding with its policy between PDPs (Policy Decision Points, like PSs) and PEPs (Policy Enforcement Points, like routers). But COPS seems to be used within intradomain networks. In this paper, the enhanced COPS is proposed to negotiate and notify domains of the policy between PDPs (for example, PSs).

3.4 Preparation

In this section, the premise behind the system structure and the preparation before starting the system are explained. Figure 8 shows a sample system structure.

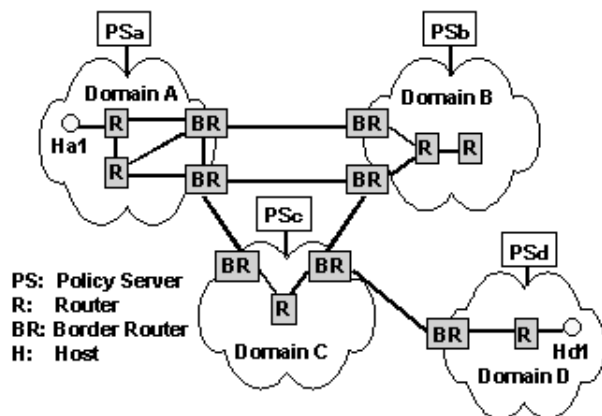


Fig. 3-2 Sample of system structure

Each PS (PSa, PSb, PSc, PSd) owns the policies of each domain. There are some routers and other devices in each domain. Border routers (BRs) in a domain are connected with other BRs in other domains. The PSs manage some devices in the domain and control them (for example, intradomain routing) according to the policy. And PSs exchange the policy among domains and execute QoS provisioning and accounting.

Network administrators must get the following information from neighboring domains, which are connected by physical communication lines, before PSs start the policy exchange process and interdomain QoS provisioning and accounting.

- Domain ID (in short, AS-ID) and IP address for PS (in short, PS-IP).
- IP address of BRs (in short, BR-IP) of the inside and outside interface.
- Resource information about the interdomain physical communication lines (which are called interdomain links, in short, I-D links) (e.g., network medium, maximum bandwidth, range of packet loss, minimum delay, minimum jitter, MTU).

Next, network administrators obtain the following information and store it in the PS's storage device.

Part of this information is treated as the policies.

- Customer information (for example, host IP, priority, available service name, right of the interdomain service use).
- AS-IDs that the above customer are allowed for interdomain communication.
- Interdomain link ID that shows that the above customer's host can arrive at the domain of the above AS-ID.

In the next section, it is assumed that PSs already have the above information.

4. Interdomain PAP

In this section, the interdomain PAP, which is an enhanced BGP4 protocol, is explained.

4.1 Outline

There are some hosts and networks that are specified by the addresses in one domain. The advertised policy is treated as a set of this information, and is distributed in another domain. A simple case is illustrated in figure 9.

The advertised policy in PSa of domain A contains at least the following information.

- Host a1 (Ha1) is in the domain A.
- To execute services using Ha1, three interdomain links (La1, La2, La3) are available that have QoS information (bandwidth, delay etc.).

This information of the policies is made by PSs that know about the structure of the domain network. For example, PSs make the flow policy of Ha1 with resource information (e.g., network medium, bandwidth) to Ha1 through La1. The policy is advertised by PAP.

Figure 9 shows how one policy in domain A is being advertised from PSa (Source PS) to PSc and PSD (Destination PS).




Fig. 4-1 Outline of PAP

This policy is advertised with routing information on BGP4. A different process is that the advertised

policy, the flow policy for example, may be transformed according to each PS in accordance with the network structure and/or state.

Firstly, PSa advertises its policy to PSc. The policy contains the following information. If a customer in domain C would like to communicate with Ha1's customer in domain A, that customer can use the 3.5-Mbit/sec bandwidth through BRc1 in domain C. The bandwidth is limited by the minimum bandwidth through the whole path between Domain A and Domain C. The available bandwidth of each link is 3.5 Mbit/sec between BRc1 and BRa2, and 10.0 Mbit/sec between BRa2 and Ha1.

Secondly, PSc forwards its policy to PSd. This policy contains the following information. If a customer in domain D would like to communicate with Ha1's customer in domain A, the customer can use the 2.7-Mbit/sec bandwidth through BRd1 in domain D. The bandwidth is limited by the minimum bandwidth through the whole path between Domain C and Domain D. The available bandwidth of each link is 4.0 Mbit/sec between BRd1 and BRc2 and 2.7 Mbit/sec between BRc2 and BRc1.

In figure 9, the bandwidth that is included in the flow policy is explained. But there are other attributes in the flow policy, which are transformed or forwarded by each PS as follows.

Attributes	Criterion
Bandwidth	Minimum value
Packet loss	Range value
Latency	Total sum value
MTU	Minimum value
Jitter	Total sum value

Thus, PS is allowed to transform the policies in this process, but this paper does not address these details.

4.2 Protocol

All PAP is based on BGP4, except for the policy transformation process.

4.3 Message format

The policies are advertised using UPDATE messages on BGP4. The format of a PAP message is shown in figure 10.

- **PAP is Extension BGP4**

- Conveying policies in UPDATE message with routing information




Fig. 4-2 Message format of PAP

5. Interdomain PNP

This section explains the PNP, which is an enhanced COPS protocol, used among domains.

5.1 Outline

The PNP is outlined and the sample policy and functions are explained.

The PNP is used to convey some policy and to execute some actions according to its policy. To control the use of these services among domains, a signaling function that triggers the control is needed to carry the necessary policy. For this PNP process, COPS can be used. Currently, COPS is only used between policy decision points like PSs and policy enforcement points like routers. Because the PNP must be enhanced to be able to adapt it for use between multiple PSs, PSs must function as both the PDP and PEP.

First, each PS makes the communication connection for the PNP. The connection process is that a PS sends a Client-Open (CO) message to another PS, and the PS that sent the CO receives a Client-Accept (CA) message from the another PS that sent a CA.

When a starting request from a QoS-guaranteed service occurs in the source domain (Domain A), the source PS (PSa) investigates several things. One is whether QoS provisioning is or is not available in the domain. After investigating, the source PS (PSa) sends a Decision (DEC(Install)) message that includes a decision policy to the destination PS (PSc).

Next, the destination PS (PSc) investigates whether QoS provisioning is or is not available in the domain, and sends a Report State (RPT (Installed/Not Installed)) message that includes its reply (accept negotiation or not) to the source PS (PSa).

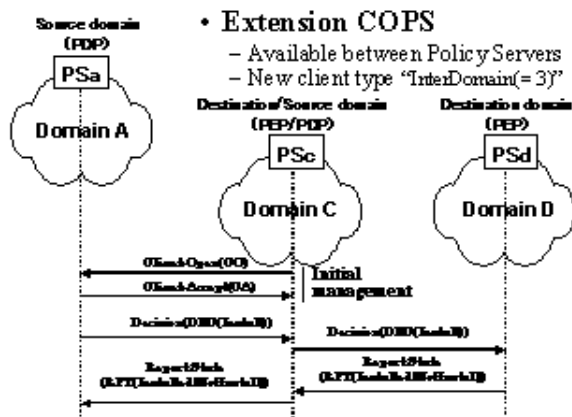


Fig. 5-1 Outline of PNP

But as figure 11 shows, there is no direct physical connection between Domain A and Domain D. In this case, both PSa and PSd need to make a logical connection through Domain C. Thus, after PSc receives the Decision (DEC (Install)) message from PSa, PSc searches for a path that can connect Domain A to Domain D, and then investigates several items about the path. After that, PSc sends another Decision (DEC (Install)) message that includes the decision policy to PSd. PSd investigates whether QoS provisioning is or is not available in the domain, and sends a Report State (RPT (Installed/Not Install)) message that includes its reply (accept negotiation or not) to PSc. After PSc receives PSd's reply, PSc sends a State (RPT(Installed/Not Install)) message to PSa.

Thus, this process is available regardless of the numbers of domains.

5.2 Protocol

All PNP is based on COPS, except for communication between PDPs (PSs).

5.3 Message format

All message formats are based on COPS.

6. Evaluation

In this section, the results of an experimental network that has PSs that use the above two protocols are presented. The experimental network was used to confirm that the interdomain QoS provisioning process is effective when service providers provide interdomain QoS services from the point of view of performance, network traffic, network administrator's tasks.

Figure 12 shows the experimental environment. Eight network administrators operate PSa and/or PSb. Each network administrator performs the same experiment four times.

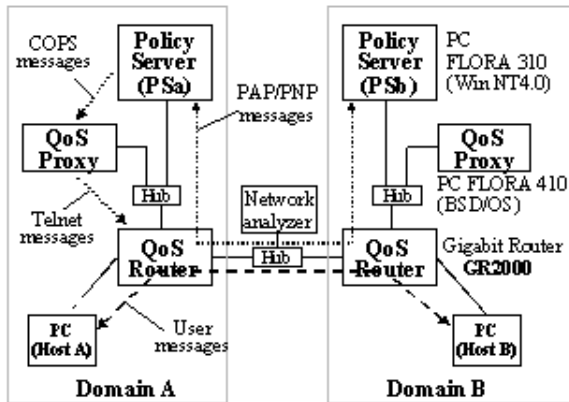


Fig. 6-1 Experimental environment

6.1 Evaluation of PAP

6.1.1 Purpose

The purpose of this experiment is to evaluate the effectiveness of PAP process from the point of view of (1) performance, (2) network traffic, and (3) network administrator's tasks.

6.1.2 Contents of experiments

1. The time taken to update the advertised policy in PS in terms of
 - transmission between PSa and PSb
 - operation of network administrator on PSa
2. The traffic on network to update the advertised policy between PSa and PSb
3. Questionnaires were also sent to network administrators. The questions focused on
 - whether or not it was easy to input the advertised policy to the PS
 - whether or not it was easy to modify the advertised policy
 - whether or not it was easy to start up PAP process.
 - whether starting up a PAP process was fast or not

6.1.3 Results and evaluations

1. The transmission time is 296.91 milliseconds on average. This shows PAP performs well in transmitting the advertised policy.
2. The operation time is 71 seconds on average. This shows that network administrators do not need more time to manage the policy advertisement.
3. The traffic is 1109.2 bytes on average. This shows that there was not so much traffic in updating the advertised policy.
4. The main information taken from the questionnaires was that network administrators felt it was hard to input the advertised policy, but they felt the policy was easy to modify. All network administrators felt that starting up and managing a PAP process was simple and fast.

6.2 Evaluation of PNP

6.2.1 Purpose

The purpose of this experiment is to evaluate the effectiveness of PNP process from the point of view of (1) performance, (2) network traffic, and (3) network administrator's tasks.

Contents of experiments

1. The time required for PNP process to succeed in terms of
 - operation for registration of the notified policy
 - operation for cancellation of the notified policy
2. The traffic on the network to convey the notified policy between PSa and PSb
3. Questionnaires were also sent to network administrators. The questions focused on
 - whether or not inputting registration of the notified policy was easy
 - whether PNP process with registration is fast or slow
 - whether or not inputting cancellation of the notified policy was easy
 - whether PNP process with cancellation is fast or slow

6.2.3 Results and evaluations

1. The operations for registration and cancellation are 15.0 seconds and 14.0 seconds on average. This shows PAP performs well in transmitting the advertised policy.
2. The traffic for operation of registration and cancellation are 1316.0 bytes and 904.0 bytes on average. This shows that there is not so much traffic during a PNP process, which should mean no problems for real network management.
3. The main information taken from the questionnaires was that network administrators felt it was easy to manage the registration and/or the cancellation of the advertised policy. However, they felt the cancellation process was slow.

The conclusion drawn from these two types of experimental results is that this interdomain QoS provisioning is effective when service providers provide interdomain QoS services.

7. Conclusion

This paper has described how to execute QoS-guaranteed provisioning services and accounting services that are offered by QoS-guaranteed services among domains. Results are as follows.

1. *Interdomain policy framework.* The paper has presented the interdomain policy framework and its application in terms of interdomain QoS provisioning and interdomain accounting. It has also presented 10 different policies. Six of these are identified as relevant for interdomain QoS provisioning and accounting control. These policies are exchanged by advertisement, request and response, and notification communications.
2. *Interdomain policy exchange architecture.* This paper has investigated architectures for policy exchange, and proposes two architectures. The architecture without SPS is better in terms of scalability, flexibility, and legacy of the current internet architectures.
3. *Interdomain policy advertisement protocol.* We proposed an interdomain PAP, which is an enhanced BGP4 protocol, for the interdomain policy advertisement.
4. *Interdomain policy negotiation and notification protocol.* We proposed a PNP, which is an enhanced COPS protocol, for the interdomain policy negotiation and notification.
5. *Evaluation.* We ran experiments to confirm the performance of interdomain QoS provisioning

process using the two protocols. The conclusion drawn from these experimental results is that this interdomain QoS provisioning is effective when service providers provide interdomain QoS services.

8. References

[INTD] T. Ebata, M. Takihiro, S. Miyake, M. Koizumi, F. Hartanto, and G. Carle, "Interdomain QoS Provisioning and Accounting", Internet-Draft, draft-ebata-interdomain-qos-acct-00.txt, November 1999.

[QOS] S. Gai, J. Strassner, D. Durham, S. Herzog, H. Mahon, and F. Reichmeyer, "QoS Policy Framework Architecture", Internet-Draft, draft-sgai-policy-framework-00.txt, February 1999.

[SLA] J. Strassner, E. Ellesson, and B. Moore, "Policy Framework LDAP Core Schema", Internet-Draft, draft-ietf-policy-core-schema-04.txt, June 1999.

[TERM] J. Strassner, and E. Ellesson, "Terminology for describing network policy and services", Internet draft, draft-strassner-policy-terms-01.txt, February 1999.

[RFC1364] K. Varadhan, "BGP OSPF Interaction", RFC 1364, September 1992.

[BGP4] Y. Rekhter, and T. Li, "A Border Gateway Protocol 4 (BGP-4)", Internet-Draft, draft-ietf-idr-bgp4-09.txt, September 1999.

[COPS] J. Boyle, R. Cohen, D. Durham, S. Herzog, R. Rajan, and A. Sastry, "The COPS (Common Open Policy Service) Protocol", Internet-Draft, draft-ietf-rap-cops-06.txt, February 1999.

