

Implementation of a Bandwidth Broker for Dynamic End-to-End Capacity Reservation over Multiple Diffserv Domains

Ibrahim Khalil and Torsten Braun

Computer Networks and Distributed Systems (RVS)
Institute of Computer Science and Applied Mathematics (IAM)
University of Berne
Neubrückstrasse 10, CH-3012 Bern, Switzerland
{ibrahim,braun}@iam.unibe.ch

Abstract. As Virtual Leased Line (VLL) type point-to-point connections over Internet are now possible with the Expedited Forwarding (EF) Per Hop Behavior (PHB), Internet Service Providers are looking for ways to sell bandwidth services to prospective corporate customers with the so called Bandwidth Brokers (BB). Based on this approach most of the recent implementations consider providing such services only in a single provider domain. However, many of the ultimate customers might actually want to extend VLLs up to the periphery of other providers. In this paper, we describe the implementation of a BB that uses simple signaling mechanism to communicate with other cooperative Brokers to enable customers to dynamically create VLLs over multiple Diffserv domains.

1 Introduction

Virtual Leased Line (VLL) type point-to-point connections can now be built over various IP segments with the recently proposed Expedited Forwarding (EF) [1] Per Hop Behavior (PHB) in the differentiated services [2] architecture. To take advantage of this new technology Bandwidth Brokers that can dynamically create VLL on demand have been proposed in [3], [4] and refined in [5], [6]. New Bandwidth Broker models based on the existing architectures have also been proposed in [7], [8], [9] and several implementations have been reported in [7], [5], [6], [10]. Most of these implementations of Bandwidth Brokers have the following characteristics in common:

- They are mostly responsible for a single Diffserv Domain. The Bandwidth Brokers are capable of advance or immediate reservation only in the domains they maintain.
- All the concepts propose policing users traffic at the ingress edge only. Except [11] most of the BBs don't consider interior provisioning.
- Almost all except [7] and [11] rely on RSVP for signaling.

While the existing Bandwidth Broker implementations don't yet have mechanisms to communicate with other neighboring domains, they mostly propose modified RSVP or similar mechanisms as the method for both inter-domain and intra-domain signaling. Although both of these have the potential to be integrated in the future advanced Bandwidth Brokers, at the moment when there are core Service Level Agreements (SLAs) and resource provisioning issues yet to be solved [12], they (i.e. the use of RSVP like signaling) might further complicate implementation issues and delay easy and rapid deployment. For example, as mentioned in [4], for resource reservation over a Diffserv network using such Bandwidth Brokers, both sending and the receiving hosts need to be present during reservation and also during the period the reserved interval starts. In reality, the sender or receiver might not even exist during the reservation process.

In this paper, keeping this in mind, we present a simple approach to make advance reservations in the absence of senders or receivers in a multi-domain scenario. Rather than using RSVP in inter-domain signaling to reserve capacity across domains, we use a novel method to identify domains, and hence the Bandwidth Brokers that are responsible for maintaining them. Section 2 presents basic components and ingredients for making reservations over several Diffserv domains with Bandwidth Brokers. Section 3 describes implementation architecture and the components in that architecture. In section 4, we describe operational details and system flows of the BB, and in section 5 we clarify the operational details by presenting some real examples. Finally, in section 6, we conclude our paper with a summary and future research directions.

2 End-to-End Capacity Reservation

2.1 An Example Scenario

Consider the scenario as shown in Figure 1. The domains are Diffserv [2] enabled and under different administrative control. This means that if stub networks C or D in domain 1 want to establish VLL with Stub network A in the same domain or with stub network B in domain 2, traffic entering domain 1 is classified and possibly conditioned at the boundaries (edge router 2) of the network, and assigned to different behavior aggregates. Each behavior aggregate is identified by a single DS codepoint (DSCP for EF). In the interior (and also egress) of the network, with the help of DSCP- PHB mapping certain amount of node resources can be allocated for this quantitative traffic.

2.2 An Automated System: Bandwidth Broker

In the example above if the administrative control of each ISP is given to an automated system like a Bandwidth Broker, its responsibilities will be :

- Check request validity. In the example, for the VLL over domains, BB 1 needs to check the validity of stub network A's request and BB 2 needs to check the request of ISP domain 1.

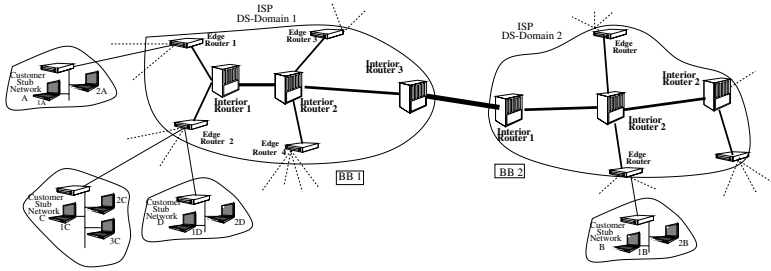


Fig. 1. Diffserv Deployment Scenario across Multiple Diffserv Domains

- Perform admission control in its domain. In a simple case, this can be only checking resource availability at the border routers as these are the obvious points that will be traversed by a VLL connection. In a more advanced case this can be checking resource availability at all the routers along the path from ingress towards egress.
- Coordinate with other separately administered Bandwidth Brokers. In the example, Bandwidth Broker 1 will need to signal to Bandwidth Broker 2 for resource reservation on behalf of stub network A. If there are several ISP domains each managed by such Brokers, the job of this coordination also means identifying the right domains and corresponding brokers for a certain resource allocation request (RAR).
- Configure ingress router of its domain if a request is accepted. Configuring ingress router means dynamically marking the traffic as EF and policing to limit the volume of traffic.

2.3 Service Level Agreements

Service Level Agreements are generally contracts between network service providers and their customers to guarantee particular quality of service levels for network performance. SLAs exist between a customer and its provider (called intra-domain or customer-ISP SLA) and also between two providers (called inter-domain or inter-ISP SLA).

A customer normally has a contract with the local ISP regarding the maximum amount of traffic he can send or receive for a VLL service. Such customer-ISP SLA, however, doesn't automatically guarantee that a customer will always receive the service upon request - it only indicates the upper limit of the request and successful reservation of the requested VLL depends on admission tests at different points along the VLL. Referring to Figure 1, if a customer wants to establish a VLL between stub network A and C, an intra-domain SLA would suffice. However, for a VLL to be established between stub network A and B and inter-domain SLA between domain 1 and 2 must be in place. Based on inter-domain SLA the sending domain can send a maximum, say X Mbps aggregated traffic, to a receiving domain, and ensures that it doesn't send more than X Mbps by shaping at the outgoing interface connected to receiving domain. Receiving

domain polices the incoming traffic at the ingress router's inbound interface to make sure that the sending domain doesn't send more than X Mbps.

2.4 Service Provisioning and Call Admission Control

Determination of resources required at each node for quantitative traffic needs the estimation of the traffic volume that will traverse each network node. While an ISP naturally knows from the SLA the amount of quantitative VLL traffic that will enter the transit network through a specific edge node, this volume cannot be estimated with exact accuracy at various interior nodes that will be traversed by VLL connections. However, if the routing topology is known, this figure can be almost accurately estimated. For simplified provisioning and admission control we assumed the following:

- Pre-configure interior and other border routers with scheduling mechanism like Priority Queuing (or CBQ, WFQ) so that traffic marked as EF are served by high priority queue.
- Traffic follows a known path.

If the domain is QoS rich, for a simple model it might suffice only to perform CAC at the edge points. For a more sophisticated model, considering the necessity of interior provisioning the BB may also check the availability of resources at the interior points that would be traversed by a VLL. In such a case, virtual core provisioning [11] might be suitable that only requires a capacity inventory of interior devices to be updated based on VLL connection acceptance, termination or modification.

2.5 End-to-End Signaling

A user sends a request to the Bandwidth Broker that maintains the user's ISP domain. The request contains source and destination addresses and also the requested bandwidth. While the source naturally resides in the stub networks attached to the ISP's network, the destination might well be in the stub network that is attached to another ISP's domain. That domain might not be the final domain and there might be one or more domains in between. If both the source and destination addresses are in the stub networks of the same ISP domain, the Broker that maintains the domain can find the ingress and egress routers by some simple lookup in the various Broker related databases (explained in the next section). The Broker performs admission control at the points (ingress and egress) before deciding whether the request should be granted or not. If the destination is in another domain other than the source domain, then the Broker must identify the followings:

- the domain that has the destination stub connected to it.
- intermediate domains (if any) that will be traversed by VLL connection if the request is accepted.

Before we investigate the above two issues it would be useful if we give a brief overview of Resource Allocations Requests.

BB Resource Allocation Request Format. A Resource Allocation Request (RAR) may come from an individual user to a BB or from one BB to another neighbor BB. We call the first one intra-domain RAR while the latter one is referred as inter-domain RAR. Their formats are:

- **Intra-domain RAR:** To setup a new VLL this request contains user id and password, source and remote addresses of the VLL to be established and the bandwidth required for it:
Newflow -u userid -p password -s source -d remote -b bandwidth
- **Inter-domain RAR:** Inter-domain RAR is automatically generated by a broker when it detects that the remote address indicated in the request is not attached to its domain. The request is then sent to another neighbor domain. Since the actual requester is a domain broker, the recipient broker needs to check its validity as an inter-domain request.
Newflow -bb brokerid -p password -s source -d remote -b bandwidth -tbb final_domain

Domain Identification. A scalable and simple way for each Broker would be to send boundaries of the domain that it maintains to other cooperative domains. By boundaries we mean the IP addresses of the edge devices. Lets consider domain 1 and 2 in Figure 2 where each of the domain is actually constituted from several edge devices. All these edge devices have unique IP addresses. If we can identify an edge router by the destination IP address in the RAR, then we can readily identify the domain, and hence the Bandwidth Broker that represents the domain. For example, when a user wants to establish a VLL to send traffic from any of the stub networks 7.7.x.x to one of the sub networks 5.5.x.x, Broker BB1 can easily identify that it has to finally communicate with BB7 by reading a domain identification database.




Fig. 2. (a) Domain 1 (b) Domain 2 (c) Domain Identification Database

Bandwidth Broker Message Forwarding. When the Bandwidth Broker identifies the Final Broker there might be one or more intermediate Brokers that need to be contacted as well for end-to-end capacity reservation. How does

Broker (first Broker or any intermediate Brokers) determine the next appropriate Broker when there are several neighbor Brokers and a VLL needs to be established over several domains? In the previous example the VLL needs to be established over domains that are managed by BB1, BB2 and BB7. If each Broker knows the neighbor brokers and by exchanging that information every Broker can build a message forwarding table as shown in Figure 3(c) and 3(d). From the table is obvious that BB2 is the intermediate broker that needs to be contacted first by sending an Inter-domain RAR from BB1 before BB2 finally sends another inter-domain RAR to BB7 on behalf of BB1.

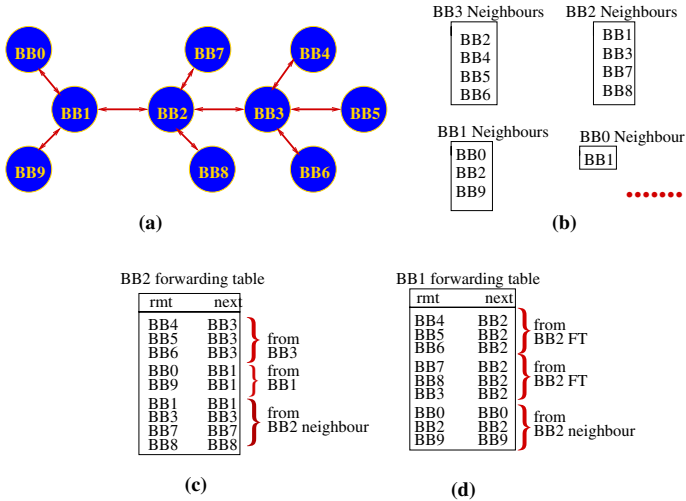


Fig. 3. (a) Several Diffserv Domains represented by Bandwidth Brokers (b) Neighbors tables of some Brokers (c) BB Message forwarding table in BB2 (d) BB Message forwarding table in BB1.

3 Implementation of the System Components of the Broker

3.1 Architecture

Based on the requirements for end-to-end capacity reservation the Bandwidth Broker has been developed to dynamically establish VLL on customer's request. Our earlier analysis and functional requirements of BB resulted in a four layer implementation architecture of Figure 4. The top layer is responsible for validating both intra- and inter-domain requests. The two middle layers are composed of several databases that are invoked for admission and signaling purposes of valid requests. The bottom layer decides and configures edge routers based on

processing of requests in the three above mentioned layers. In the next few sections we will describe the components of these layers.

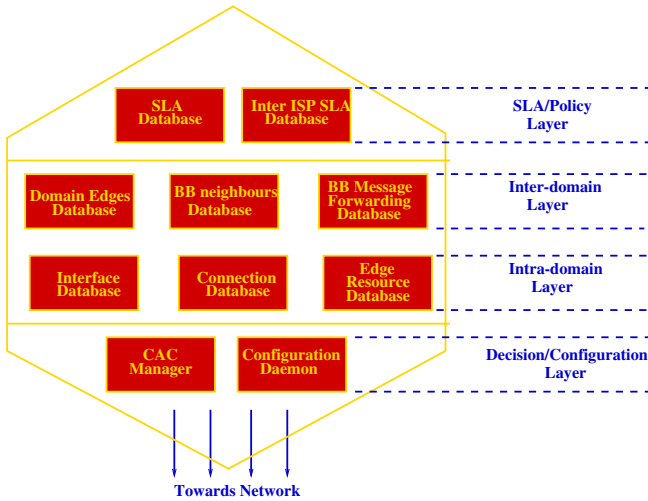


Fig. 4. Layered Implementation Architecture: Components of a BB for Resource Provisioning, Admission Decision and end-to-end Signaling

3.2 The Databases of the Bandwidth Broker

The **customer-ISP SLA database** contains not only the user's identification, but also specifies the maximum amount of traffic one can send and/or receive over a VLL. As VLL might be established between two points (i.e. source and destination) out of several possibilities, a SLA also contains the boundary of a valid VLL area and are put in this database as source and remote stub addresses. User authentication process prohibits malicious users to setup unauthorized VLL and access network resources illegally. It contains the following tuple:

<User ID, Password, Maximum BW in Mbps, Source Stub Address, Remote Stub Address>

The **inter-ISP SLA database** is invoked by a domain when it receives inter-domain RAR. By doing so the receiving domain can check the validity of the request sent by the sending domain. Here this validity means identification of the sending domain and the maximum amount of bandwidth it can reserve on a certain edge interface in the receiving domain that is directly connected to it.

<Domain ID, Domain Password, Maximum BW>

The **interface database** contains necessary records of edge routers that are used as VLL end-points for the outsourced VLL model. In such a model since some customer stub networks are connected to the ISP edge router we need to

specify which stub networks are connected to a particular edge router. Also, an edge router might have one or more inbound and out-bound interfaces which also need to be specified for each stub network that is actually connected to a particular in-bound interface of a router. This is important because normally at the in-bound interface VLLs are policed on individual basis and at the out-bound they are shaped on an aggregated basis. The tuples are :

< stub network, edge router, generic router name, in-bound interface, out-bound interface >

The **connection database** contains a list of currently active VLLs. When a request for a new VLL connection or termination of an existing connection arrives, the BB can check if that connection already exists or not and then make its decision. The storage of detail connections indicates the amount of resources consumed by VLL users at various edge and interior nodes.

< user id, source address, VLL ID, rmt address, bandwidth, activation time >

The **edge resource database** contains information regarding resource provisioned (C_{TOTAL}) for different router interfaces and used (allocated) capacity ($C_{allocated}$) to existing VLL connections. The difference between the two is the spared capacity that can be allocated to incoming connections. The tuples are, therefore:

< edge router, C_{TOTAL} , $C_{allocated}$ >

The **VLL ID database** maintains a list unique VLL IDs and their status for each edge router. An ID that is available is marked as 1, and the one that is used is marked as 2. The tuples are: *< edge router, Tunnel ID, Status >*

The **BB Neighbor Database** hold records of neighbor Bandwidth Brokers IP addresses as well as IP address of the router interfaces (both in-bound and out-bound) that interconnect the peer domains.

< Neighbor BB, InsideInterface, OutsideInterface >

The **Domain Edges (or Identification) Database** hold records of the networks that reside at the periphery of a domain. Its purpose is described in details in the previous section . An example entry of this database is also shown in Figure 2.

< Stub Network, BB >

The **BB Message Forwarding Database** contains next hop BB's IP address to send resource allocation request to final Remote Bandwidth Broker.

< Remote BB, NextBB >

3.3 CAC Manager and Configurations Daemons

CAC manager is a functional engine of BB that basically invokes the databases described above and decide the fate of an incoming request by performing admission control at various network nodes. *Configuration Daemons* are intelligent provisioning agents that are able to translate user request and BB generated pseudo rules into device specific rules to configure the routers/switches since we might have several different devices from various vendors.

3.4 Inter-domain Signaling

From earlier sections we have seen that a Bandwidth Broker not only receives RARs from a customer of its own domain or other BBs, but also sends RAR to neighbor BBs. Therefore, we have designed a Bandwidth Broker that consists of server and a client Socket program. When a Broker's Server receives a request from a client and finds itself to be the final destination BB it can convey the CAC decision back to the client, otherwise the Server tells the client of that Broker to talk to the appropriate neighbor Broker's server. So, there is a chain of communications which are handled by client-server concatenations. This is shown in Figure 5.




Fig. 5. Client-Server Concatenation for Inter-domain Signaling

4 Operational Details and System Flows

In this section we will describe how a connection is established or terminated, how various components interact with the BB, and under which circumstances a new connection request or termination request gets refused.

4.1 VLL Establishment

Figures 6 and 7 show all the communications involved in setting up a VLL connection between two stub networks or simply between an originating host and a remote host. Both intra and inter-domain cases are explained. Although an intra-domain scenario is not the focus in this paper, yet we describe it because its similarities in system dynamics with an inter-domain case, and many of the communications involved in an intra-domain scenario are actually repeated in the latter one. We will start describing the operational details by referring to the communications marked on Figure 6. Considering each communication in turn :

- 1) A user sends a VLL connection request message to the BB via http or other interfaces able to communicate to the BB server.

- 2,3) The BB contacts the customer SLA database that is responsible for validating the user and his request. If the user is identified correctly, his source and remote address conforms the contract, and also the bandwidth requested is less than or equal to the agreed traffic contract, it proceeds further.
- 4,5) The BB contacts the configuration daemon to check its status. The status can be busy, available, or down. Only in the case of availability the user request can be processed further.
- 6,7) The BB contacts the connection database to check the existence of an exactly similar VLL. This is because for a source and destination pair only one VLL can remain active.
- 8,9) The BB reads domain edges database to find out whether the VLL is needed to be created only in the domain under its supervision or might well span over other autonomous domains.

Intra-domain Case. If (Figure 6(a)) the BB finds that both source and destination are in the same domain,i.e. the VLL is needed to be created over a single domain, it proceeds as follows :

- 10,11) BB reads the interface database to find out ingress and egress edge routers. One or both are configured depending on a traffic contract.
- 12,13) Once the edge routers are detected from the interface database the BB communicates with the resource database and performs admission control on certain router interfaces to allocate a VLL of the requested amount. It might perform admission control on only the appropriate edge router interfaces or even on the interior routers interfaces that can be detected from the topology database. The resource database responds to the BB and either allocates the resource or denies based on resource availability.
- 14) The BB tells the configuration daemon to create appropriate configuration scripts. This is to be noted that configuration script is created only for the ingress edge router because this the only router that is configured to mark and police the incoming traffic. In the case of double edged SLA the egress router is configured as well for allocating QoS in the other direction. In the meantime, the resource and the connection database update their records. Another point to note is that BB also fetches a VLL ID from the VLL ID database that is unique for each VLL and needed while configuring the router. The new connection request data is appended to the connection database and the VLL ID that has just been allocated from the VLL ID database is marked as used.
- 15,16) The CD puts a busy signal on itself and creates the routing scripts. It then sends configuration scripts to the routers. The routers send signals to the CD.
- 17,18) The CD removes the busy signal from itself and sends acknowledgment to BB which sends a notification to the user.

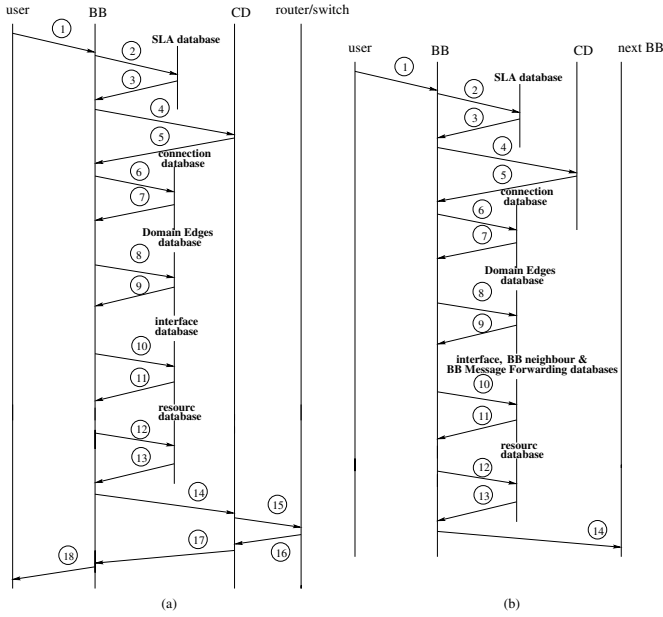


Fig. 6. VLL Establishment System Flow: (a) Intra-domain Case, (b) Inter-domain Case

Inter-domain Case. In the case a VLL (Figure 6(b)) is supposed to be established over several DiffServ domains the BB follows the steps described below:

- 10,11) Once the final destination domain has been determined the Bandwidth Broker finds out the next hop BB by reading the BB Message Forwarding Database. Now a search in the BB Neighbor Database gives current domain's outgoing interface towards the next hop BB. The BB also fetches the appropriate ingress router interface from the interface database.
- 12,13) These steps are similar to the steps 12 & 13 in the previous case. As the BB now knows the ingress and egress router interfaces, it performs admission control on those interfaces.
- 14) A positive CAC response leads to sending an inter-domain RAR to the next hop BB.

Next Hop BB as Final BB

If the next hop BB finds that the destination stub is in the domain maintained by it, the following steps are followed (Figure 7(a)):

- 15,16) Upon receiving an inter-domain RAR the next hop BB contacts inter-ISP SLA database to check the validity of the request.
- 17,18) It reads the BB Neighbor and interface databases to identify ingress and egress interfaces.
- 19,20) BB contacts the resource database and performs admission control on the previously identified ingress and egress interfaces.

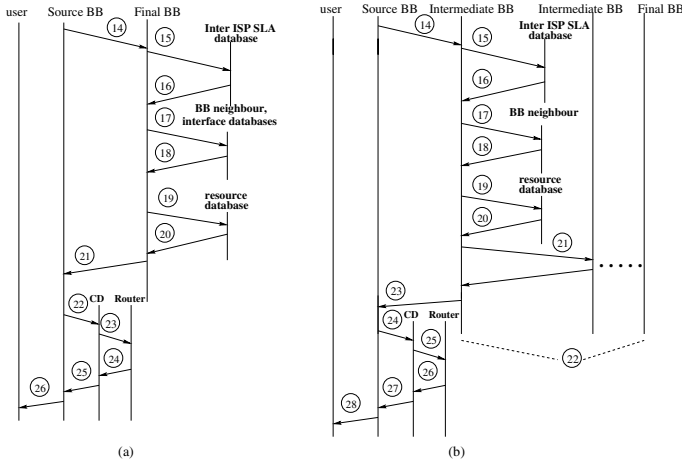


Fig. 7. VLL Establishment System Flow:Inter-domain Case cont'd. (a) Next Hop BB as Final BB (b) Next Hop BB as Intermediate BB

- 21) The BB sends CAC decision to the sender BB.
- 22-26) If the response received by the sender BB is positive then it contacts the appropriate configuration daemon to configure the ingress edge router. These steps are essentially the same like what we have seen in steps 14-18 in the intra-domain case.

Next Hop BB as Intermediate BB

The behavior of an intermediate BB is similar to that of a final BB with the exception that this one generates an inter-domain RAR based on positive CAC response from the resource database . The RAR is sent to the next BB which might be another intermediate BB or a final one. Figure 7(b) illustrates this case.

4.2 VLL Termination and VLL Request Rejection

VLL termination process involves the followings:

- The VLL connection entry is deleted from the connection database of the origin domain. Only the ingress edge router is configured to reflect the connection release.
- The resource databases are updated in all the domains that are traversed by the VLL, i.e. as resources are released $C_{allocated}$ is update as $C_{allocated} + C_{vll}$ where C_{vll} is the capacity of terminated connection.

A VLL request is rejected if

- user's SLA profile doesn't match in the origin domain, or in the case when inter-domain RAR is sent from one domain to the next neighbor domain,

interISP SLA profile of ISP that sends RAR doesn't match in the received domain.

- VLL connection already exists in the connection database of the origin domain.
- Admission control fails in any of the domains that are traversed by the VLL.

5 Examples of Dynamic Admission Control and Configuration with a BB

To test our implementation of the Broker System and its capabilities to setup VLL we ran some experiments over the public SWITCH [13] network between Bern and Geneva. The topology we used is shown in Figure 8. We have two domains with several end-systems that have private addresses and all these machines are connected to routers having public IP addresses. The domain in Geneva is represented by Broker 130.92.65.29 and the domain in Bern is managed by Broker 130.92.65.40. We also statically created VPN tunnels between these private stub networks so as to allow transparent connections between them. A Bandwidth Broker is only expected to dynamically configure ingress edge router assuming that the routers along the way from source to destination have been pre-configured with CBQ or WFQ.




Fig. 8. Experimental Setup for Demonstration of Dynamic VLL Creation over Multiple Domains

Based on the setup as shown in Figure 8 we will now explain when a VLL is established over several Diffserv domains (Figure 9, 10). Assume that user *ibrahim* plans send traffic from 172.17.0.103 to 172.18.0.100. The broker 130.92.65.40 receives the request as: `newflow -u ibrahim -p ***** -s 172.17.0.103 -d 172.18.0.100 -b 3`. As the Broker realizes that 172.18.0.0 is in domain Geneva it performs admission control in its domain (i.e. at 130.92.66.1) and then send an inter-domain RAR to 130.92.65.29 in form of `newflow -bb 130.92.66.29 -p ***** -s 172.17.0.103 -d 172.18.0.100 -b 3 -tbb 130.92.65.29`. When the broker 130.92.65.29 receives this request it knows that the request has come from another neighbor Broker (because of the tagging -bb) and therefore checks the interISP SLA

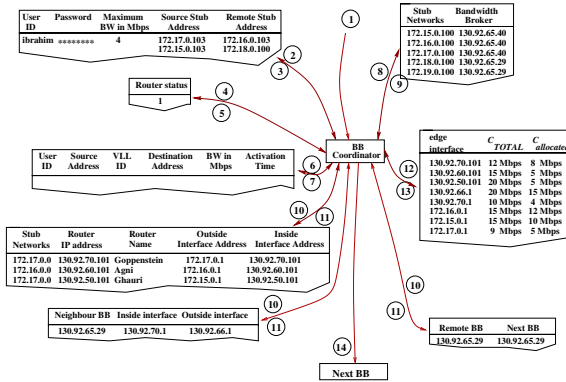


Fig. 9. An example of VLL Setup in Multi-Domain Scenario.

database to check the validity of the request. The Broker also identifies that 172.18.0.0 is located in the stub network attached to its domains since this is the final domain (from -tbb 130.92.65.29). While following the steps as described in the previous section it identifies the ingress and egress router interfaces to be 129.194.8.2 and 172.18.0.1, performs admission control on those and finally conveys the decision to the sender Broker 130.92.65.40. Upon receiving the decision Broker 130.92.65.40 talks to VLL ID database to pick up an ID, configures the edge router 130.92.70.101, and then conveys acknowledgment to the user.

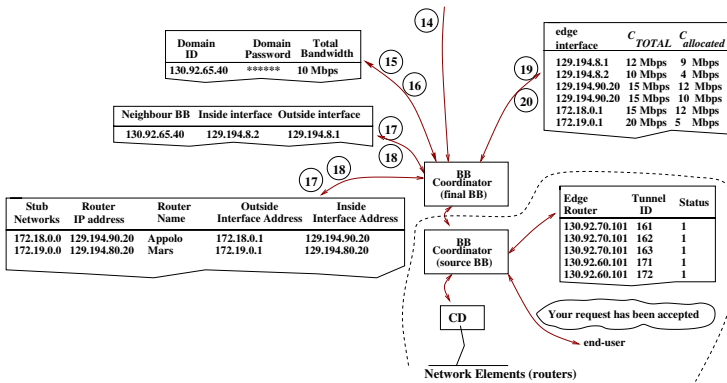


Fig. 10. Multi-domain VLL Setup Example Cont'd from Figure of 9

6 Conclusion

In this paper we have described the implementation of a Bandwidth Broker that uses a simple signaling mechanism to communicate with other cooperative

Brokers to enable customers to dynamically create VLLs over multiple Diffserv domains. We have presented a simple approach to make advance reservations in the absence of senders or receivers in a multi-domain scenario. Rather than using RSVP or COPS in inter-domain signaling to reserve capacity across domains, we used a novel method to identify domains, and hence Bandwidth Brokers that are responsible for maintaining them. A detailed implementation of the system and its operational details and some practical examples show how a simple resource reservation can be made dynamically over several cooperative Diffserv capable domains. Further simulation work might be useful to examine the scalability and effectiveness of our approach and is a topic of future research.

References

1. V. Jacobson, K. Nichols, and K. Poduri. An Expedited Forwarding phb, June 1999. RFC 2598.
2. S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weis. An Architecture for Differentiated Services, December 1998. RFC 2475.
3. K. Nichols, Van Jacobson, and L. Zhang. A Two-bit Differentiated Services Architecture for the Internet, July 1999. RFC 2638.
4. Olov Schelen. *Quality of Service Agnets in the Internet*. PhD thesis, Lulea University of Technology, August 1998.
5. Andreas Terzis, Lan Wang, Jun Ogawa, and Lixia Zhang. A Two-Tier Resource Management Model for the Internet. In *IEEE Global Internet'99*, December 1999.
6. Benjamin Teitelbaum and et al. Internet2 QBone: Building a Testbed for Differentiated Services. *IEEE Network*, September/October 1999.
7. Ibrahim Khalil, Torsten Braun, and M. Günter. Implementation of a Service Broker for Management of QoS enabled VPNs. In *IEEE Workshop on IP-oriented Operations & Management (IPOM'2000)*, September 2000.
8. Hemann De Meer, Aurelio la Corte, Antonio Puliafito, and Orazio Tomarchio. Programmable Agents for Flexible QoS Management in IP Networks. *IEEE Journal on Selected Areas in Communications*, 18(2), February 2000.
9. Zhi-Li Zhang, Zhenhai Duan, Lixin Gao, and Yiwei Thomas Hou. Decoupling qos control from core routers: A novel bandwidth broker architecture for scalable support of guaranteed services. *ACM SIGCOMM 2000*, August 2000.
10. Ibrahim Khalil and Torsten Braun. Edge Provisioning and Fairness in DiffServ-VPNs. *IEEE International Conference on Computer Communication and Network (I3CN)*, Oct 16-18 2000.
11. I. Khalil and T. Braun. Implementation of a Bandwidth Broker for Dynamic End-to-End Resource Reservation in Outsourced Virtual Private Networks. *The 25th Annual IEEE Conference on Local Computer Networks (LCN)*, November 9-10 2000.
12. Y. Bernet, J. Binder, M. Carlson, B. E. Carpenter, S. Keshav, E. Davies, B. Ohlman, D. Verma, Z. Wang, and W. Weiss. A Framework for Differentiated Services. Internet Draft `draft-ietf-diffserv-framework-02.txt`, February 1999. work in progress.
13. SWITCH. The switchlan backbone. <http://www.switch.ch/lan/national.html>.