

Bandwidth Reservation

William Johnston and Gary Hoo, Lawrence Berkeley National Laboratory

We propose a model for bandwidth reservation that can be used in the context of a general resource reservation scheme, but at the same time stay within the scalable model of the differentiated classes of service as described in the IETF diffserv Working Group documents ([5]).

The basic idea is to have bilateral end node agreements that “reserve” bandwidth in the sense that a site actively manages allocation against one or more classes of service. The overall limits on a class of service are established in the corresponding service-level agreement between the institution of the end nodes and the ISP, but the allocation of flows to this class is closely managed by the end node institutions at the site egress.

Further, the resource allocation should be policy based in a way that allows automated reservation, and it should also be possible to proxy one’s policy based authority to another site so that the bilateral agreements necessary for inter-site application operation happen automatically. (See, e.g., [2].)

The network level technology to accomplish this is provided by the classifier/shaper/policer functions of the diffserv “traffic conditioner” (TC) element. Layered on top of the TC is a “slot” allocation mechanism (“bandwidth manager”) that manages the use of a service class. When instantiated, this slot is a “micro flow” in the diffserv terminology.

Reservation requests are made to the bandwidth manager. The identity of the requestor (user_A), together with the requested resource (time slot, source id, bandwidth) are compared with policy. If the requestor and resource meet the policy, a reservation is made (the slot is allocated and the available bandwidth in the SLA is decremented) and a certificate (a digitally signed document) is issued by the bandwidth manager to represent the reservation.

When, at some point in the future, a request is made to instantiate the flow (i.e. start the instrument or application) the bandwidth manager retrieves the certificate (based on the requestor id and flow characteristics), validates the user and certificate, and instantiates the flow.

The flow characteristics are passed to the TC for classification and enforcement. From the point of view of the ingress router of the ISP, the SLA is never violated because the site bandwidth manager does not over allocate and the TC enforces flow characteristics as reserved.

Another important component in this architecture is a bandwidth broker. This service interacts with the bandwidth manager at the target site in order to accomplish the bilateral reservation. The model here is that some entity (“user_B”) at the target site (“site_B”) is the (willing) receiver of the flow. The site_B entity must have the right (i.e., be within the policy of site_B) to utilize this flow. User_B conveys (a priori) its authority (in the form of a proxy certificate) to user_A, and the site_A bandwidth broker presents this proxy to the site_B bandwidth manager in order to accomplish the reservation. The site_B incoming flow could probably just be authenticated based on the flow spec matching the reservation (i.e., site_B trusts site_A to authenticate the flow when it is instantiated), although more elaborate authentication is possible.

Akenti Application: Bandwidth Reservation

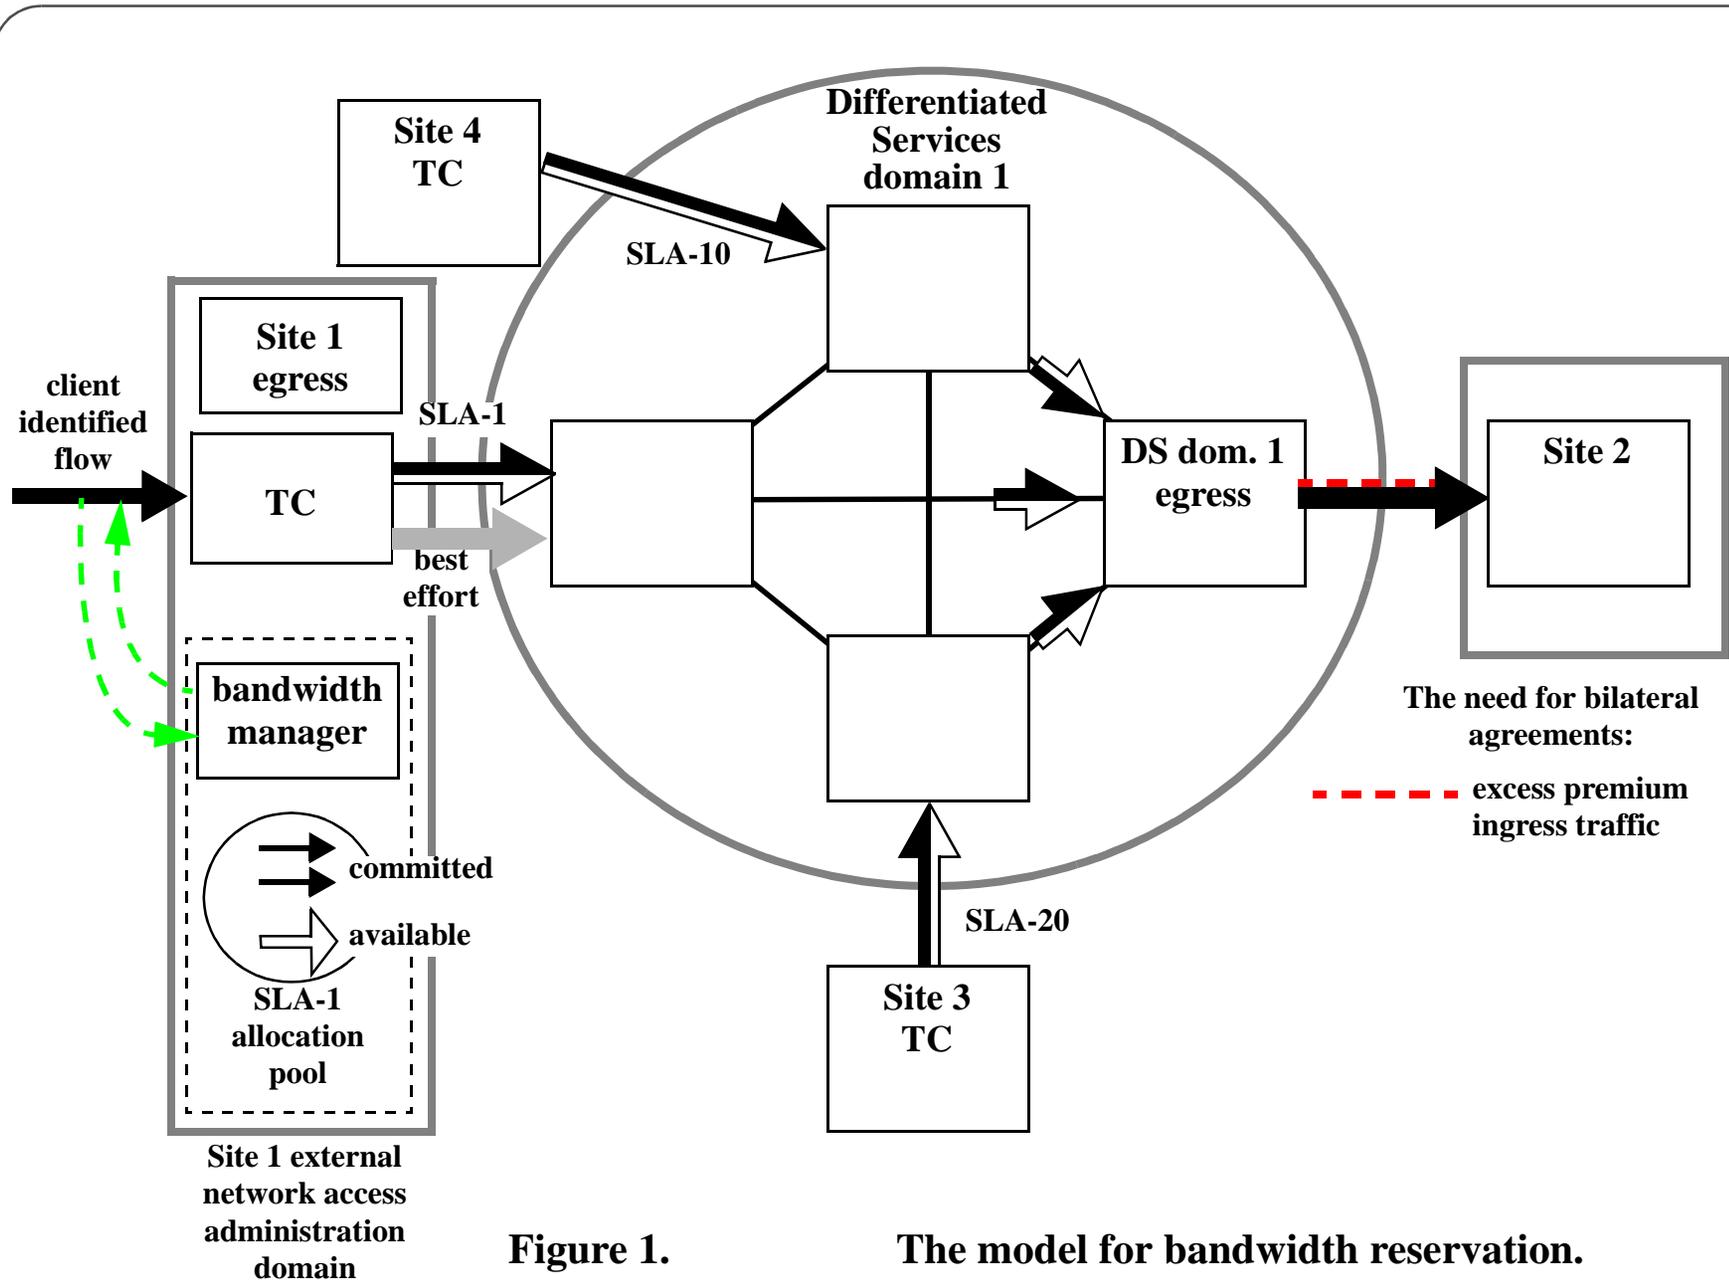


Figure 1.

The model for bandwidth reservation.

Akenti Application: Bandwidth Reservation

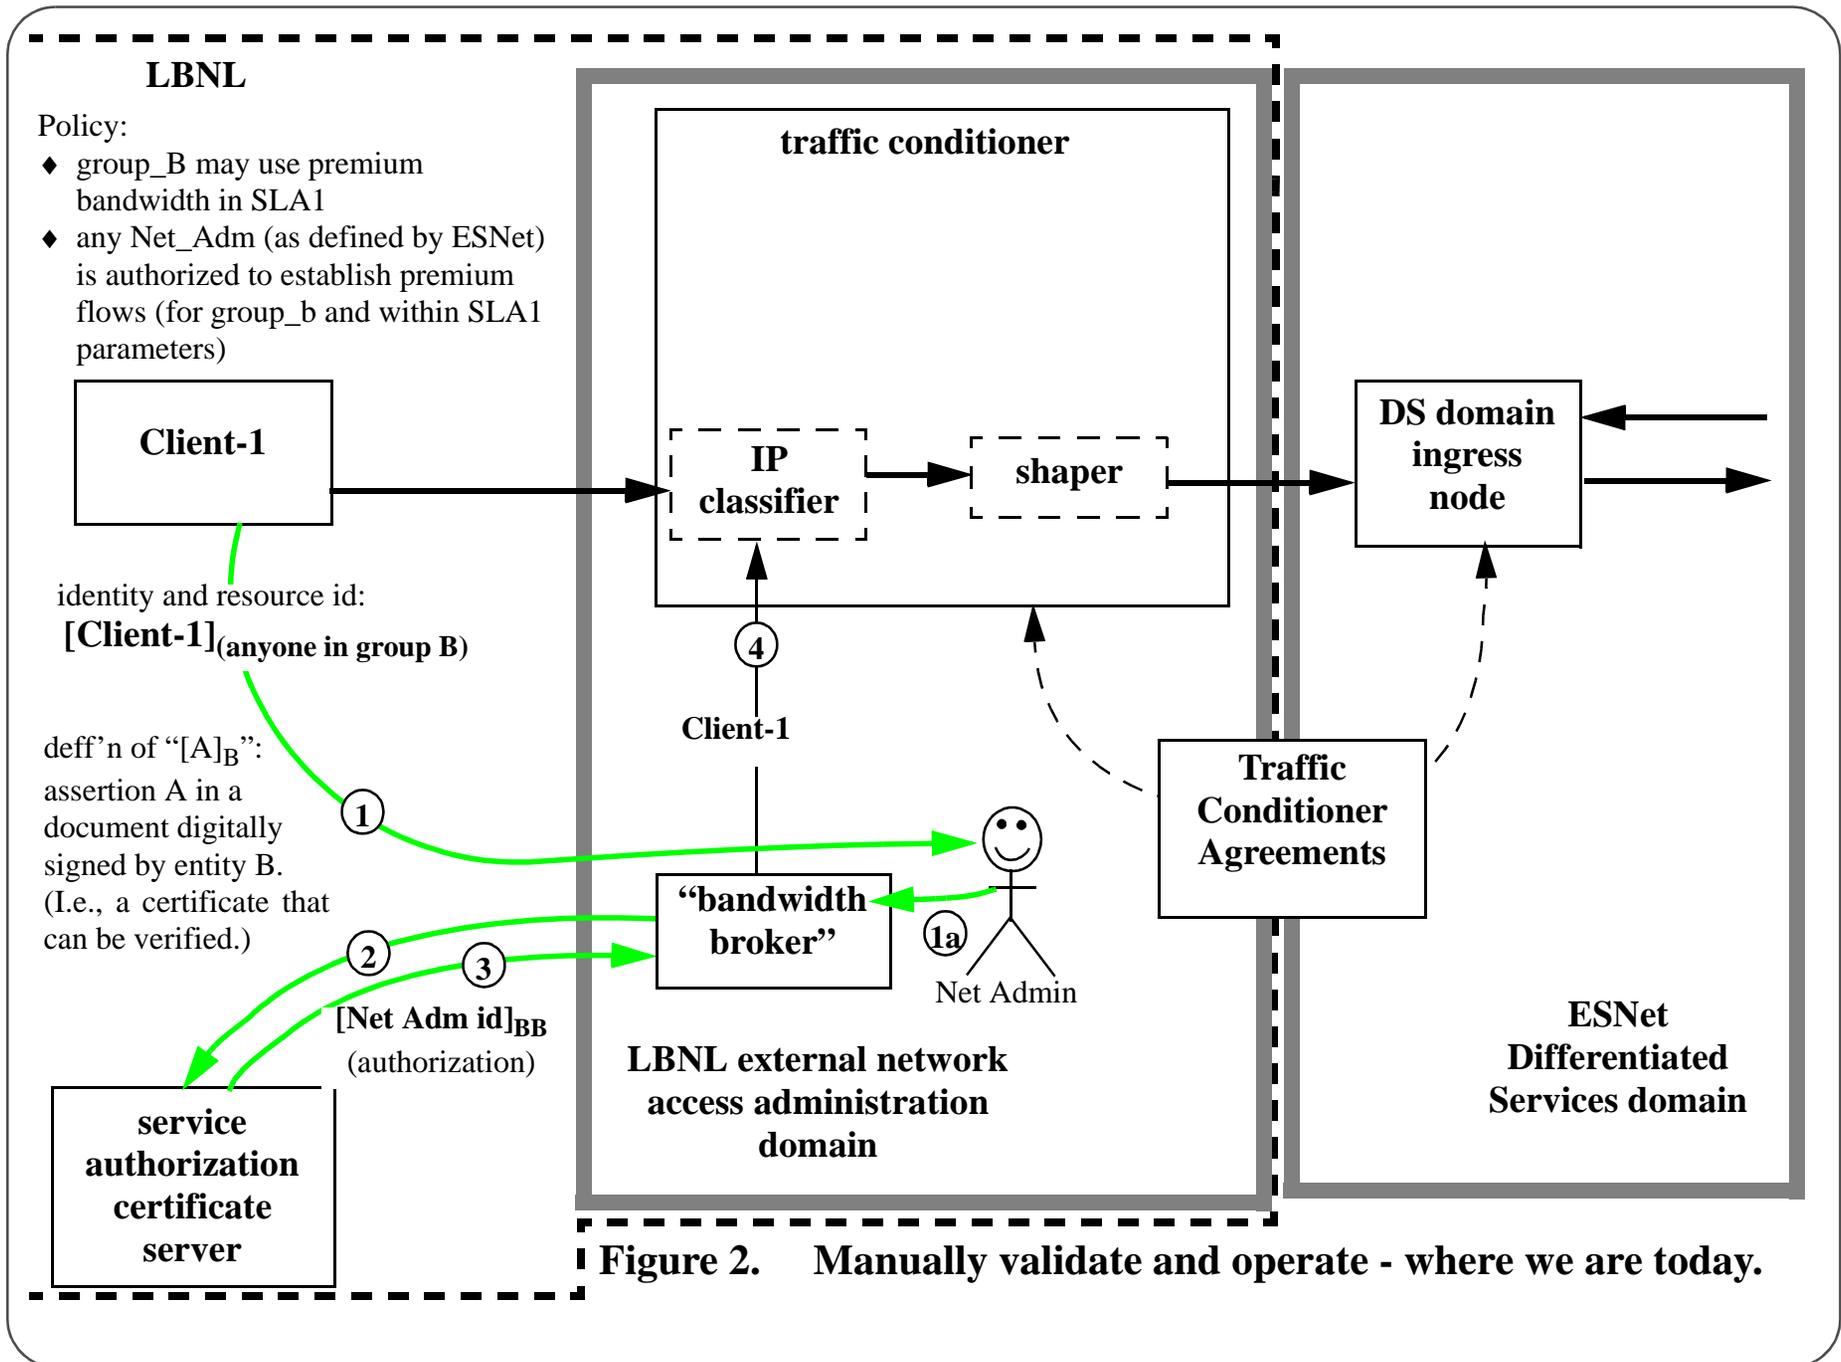
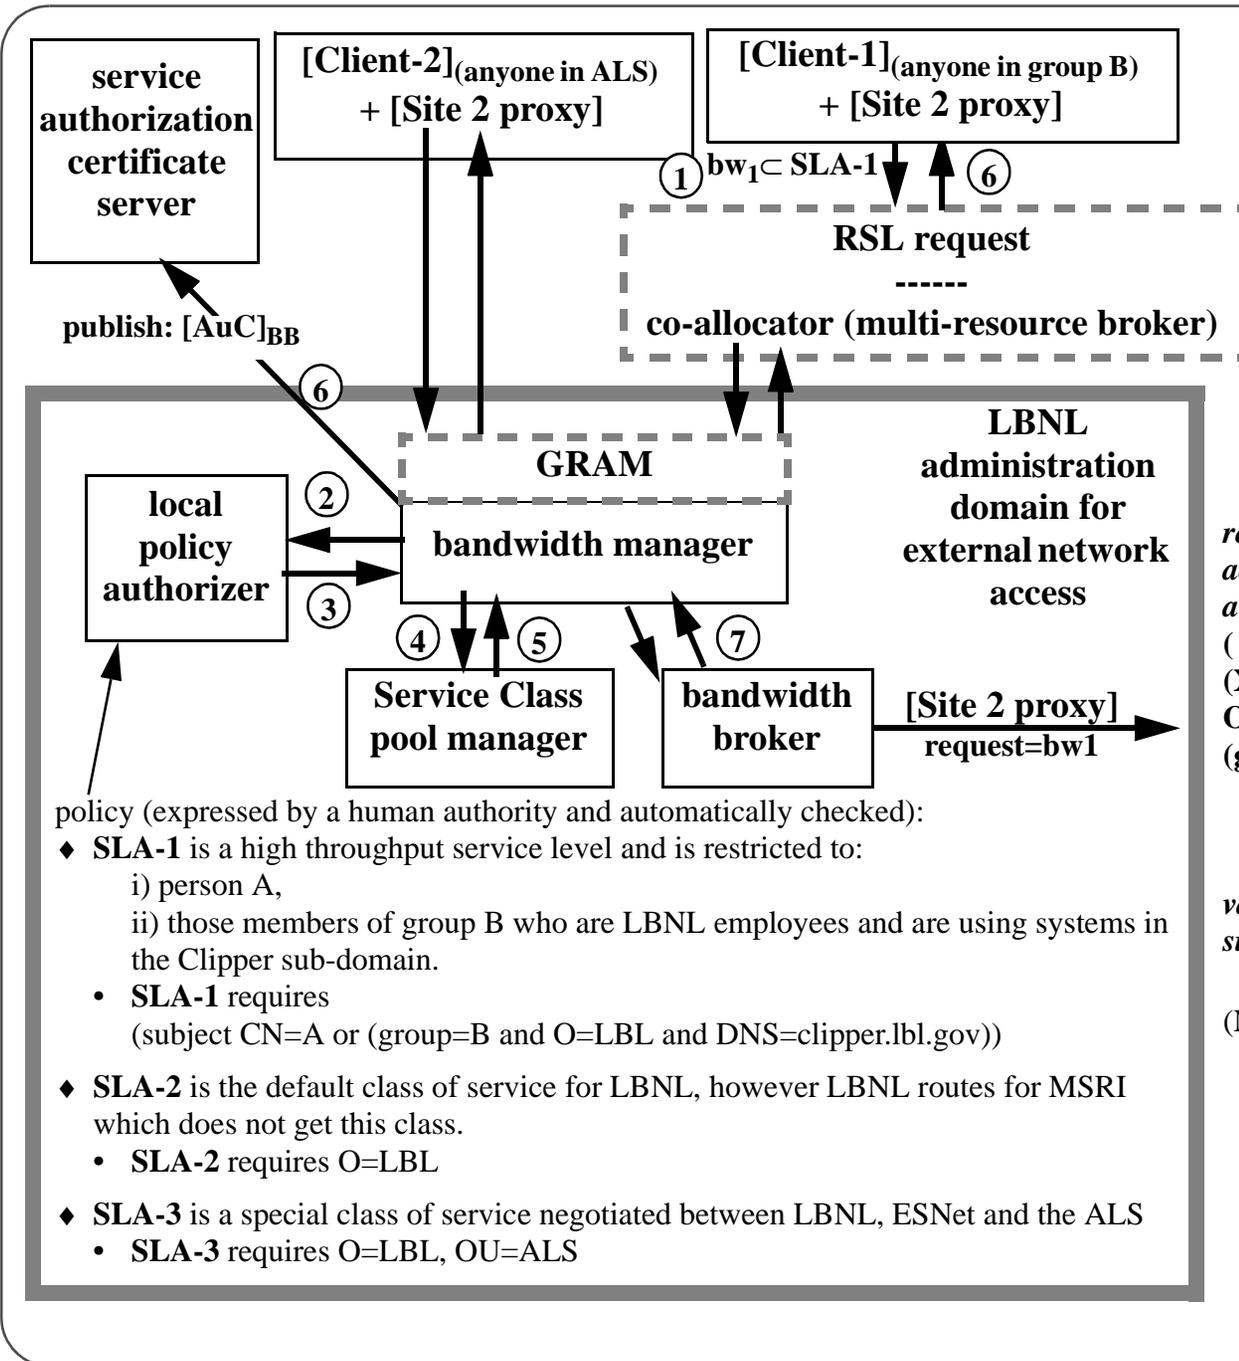


Figure 2. Manually validate and operate - where we are today.

Akenti Application: Bandwidth Reservation

Figure 3. Premium bandwidth request, policy-based authorization, and reservation phase (automated).



Authorization certificate ("AuC")
 (the result of a policy-based decision to permit the use of this class of service)

resource = SLA-1
action = 100 megabit/sec
authorized =
 (
 (X.509@LBL-ca.lbl.gov:CN=A)
 OR
 (group=B@domain_A and
 X.509@LBL-ca.lbl.gov:O=lbl.gov
 AND DNS=clipper.lbl.gov)
)
validity times = hour1:day1 - hour2:day2
signed = [message digest]_{BB}

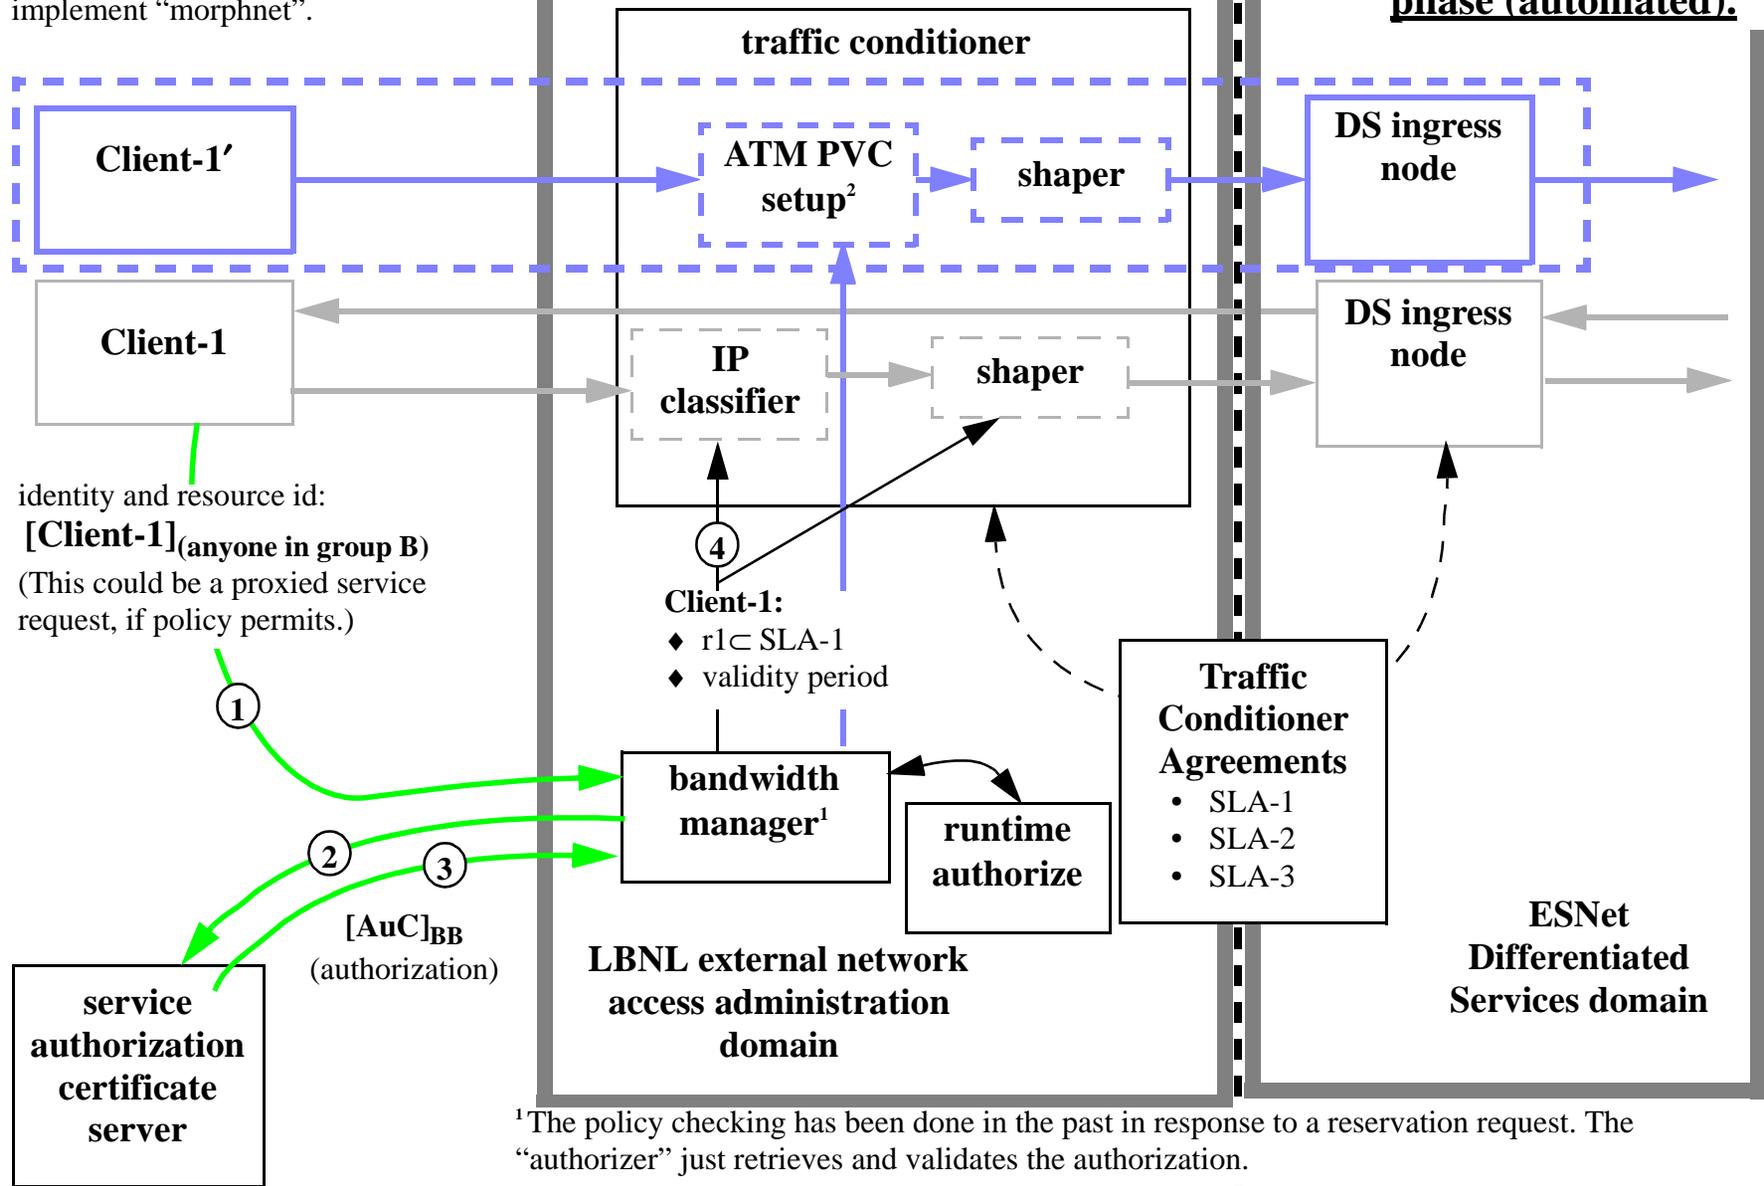
(Note that "DNS=" is a runtime check.)

⇒ **Globus resource allocation mechanisms**

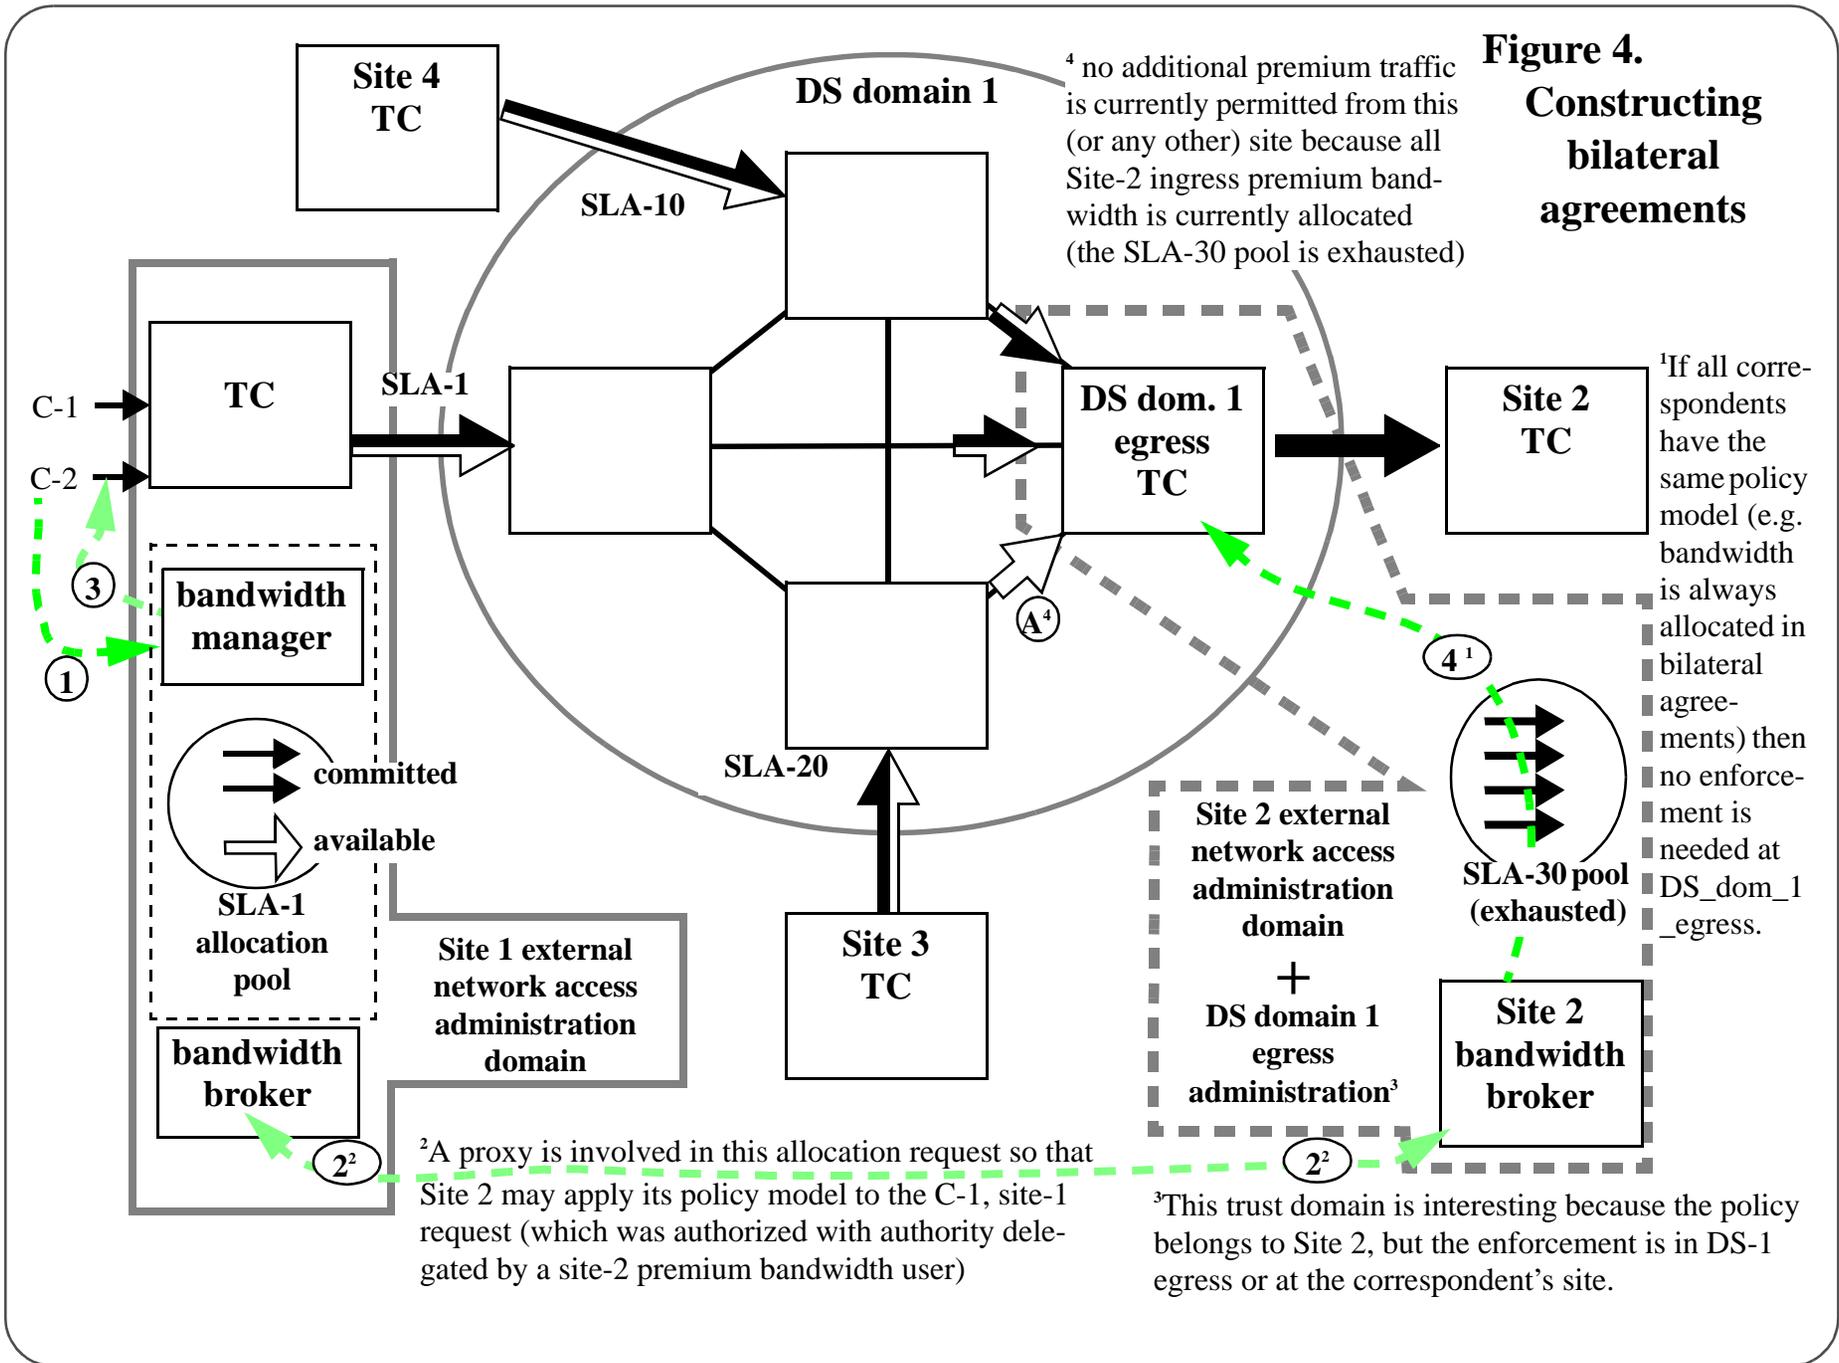
Akenti Application: Bandwidth Reservation

²A very similar model could be used to implement “morphnet”.

Validate and operate phase (automated).



Akenti Application: Bandwidth Reservation



Akenti Application: Bandwidth Reservation

Notes

Unless otherwise noted, these paper are on the Web, and pointers may be found at <http://www-itg.lbl.gov/~johnston/papers> .

[1] **“The NetLogger Methodology for High Performance Distributed Systems Performance Analysis,”**

Brian Tierney, W. Johnston, J. Lee, G. Hoo, C. Brooks, D. Gunter. 7th IEEE Symposium on High Performance Distributed Computing, Chicago, Ill. July 29-31, 1998.

[2] **“Authorization and Attribute Certificates for Widely Distributed Access Control,”**

William Johnston, S. Mudumbai, and M. Thompson. IEEE 7th International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises - WETICE'98, Stanford, CA. June, 1998.

[3] **“Real-Time Generation and Cataloguing of Large Data-Objects in Widely Distributed Environments,”**

W.Johnston, Jin G., C. Larsen, J. Lee, G. Hoo, M. Thompson, and B. Tierney (LBNL) and J. Terdiman (Kaiser Permanente Division of Research). Invited paper, International Journal of Digital Libraries - Special Issue on “Digital Libraries in Medicine”. May, 1998.

[4] **Akenti**

PKI, Attribute, and Use-Condition certificate based access control with distributed management of multi-party policy.
See <http://www-itg.lbl.gov/security/Akenti>

[5] **diffserv**

“There is a clear need for relatively simple and coarse methods of providing differentiated classes of service for Internet traffic, to support various types of applications, and specific business requirements. The differentiated services approach to providing quality of service in networks employs a small, well- defined set of building blocks from which a variety of services may be built.”
<http://www.ietf.org/html.charters/diffserv-charter.html>

[6] **WALDO**

“The Wide Area Large Data Object Architecture: We are exploring the use of highly distributed computing and storage architectures to provide all aspects of collecting, storing, analyzing, and accessing large data-objects. These data-objects can be anywhere from tens of MBytes to tens of GBytes in size. They are typically the result of a single operational cycle of an instrument, such as: single large images from electron microscopes, video images from cardio-angiography, sets of related images from MRI procedures and images and

Akenti Application: Bandwidth Reservation

numerical from a particle accelerator experiment. The source of such data objects, e.g. centralized health care facilities or large scientific instruments is often remote from the users of the data and from available large-scale storage and computation systems.”
“Our Large Data-object Architecture utilizes a high-speed wide-area ATM network between the object sources and a mutli-level distributed storage system (DPSS). As the data is being stored, a cataloguing system (ImgLib) automatically creates and stores condensed versions of the data, textual metadata and pointers to the original data. The catalogue system provides a Web based graphical interface to the data. The user is able the view the low-resolution data with a standard internet connection and Web browser, or if high-resolution is required can use a high-speed connection and special application programs to view the high-resolution original data.”
See <http://www-itg.lbl.gov/WALDO>

**This document, and more information, may be found at
<http://www-itg.lbl.gov/security/Akenti>**