

Network Working Group
Request for Comments: 2905
Category: Informational

J. Vollbrecht
Interlink Networks, Inc.
P. Calhoun
Sun Microsystems, Inc.
S. Farrell
Baltimore Technologies
L. Gommans
Enterasys Networks EMEA
G. Gross
Lucent Technologies
B. de Bruijn
Interpay Nederland B.V.
C. de Laat
Utrecht University
M. Holdrege
ipVerse
D. Spence
Interlink Networks, Inc.
August 2000

AAA Authorization Application Examples

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2000). All Rights Reserved.

Abstract

This memo describes several examples of applications requiring authorization. Each application is described in terms of a consistent framework, and specific authorization requirements of each application are given. This material was not contributed by the working groups responsible for the applications and should not be considered prescriptive for how the applications will meet their authorization needs. Rather the intent is to explore the fundamental needs of a variety of different applications with the view of compiling a set of requirements that an authorization protocol will need to meet in order to be generally useful.

Table of Contents

1. Introduction	3
2. PPP Dialin with Roaming	4
2.1. Descriptive Model	4
2.2. Authorization Requirements	6
3. Mobile-IP	6
3.1. Relationship to the Framework	10
3.2. Minimized Internet Traversal	10
3.3. Key Distribution	10
3.4. Mobile-IP Authorization Requirements	11
4. Bandwidth Broker	12
4.1. Model Description	13
4.2. Components of the Two-Tier Model	13
4.3. Identification of Contractual Relationships	13
4.3.1. Single-Domain Case	14
4.3.2. Multi-Domain Case	15
4.4. Identification of Trust Relationships	16
4.5. Communication Models and Trust Relationships	18
4.6. Bandwidth Broker Communication Models	19
4.6.1. Concepts	19
4.6.1.1. Intra-Domain Authorization	19
4.6.1.2. Inter-Domain Authorization	19
4.6.2. Bandwidth Broker Work Phases	20
4.6.3. Inter-Domain Signaling	20
4.6.3.1. Phase 0	20
4.6.3.2. Phase 1	20
4.6.4. Bandwidth Broker Communication Architecture	22
4.6.5. Two-Tier Inter-Domain Model	23
4.6.5.1. Session Initialization	23
4.6.5.2. Service Setup	23
4.6.5.3. Service Cancellation	24
4.6.5.4. Service Renegotiation	24
4.6.5.5. RAR and RAA	24
4.6.5.6. Session Maintenance	24
4.6.5.7. Intra-domain Interface Protocol	24
4.7. Requirements	24
5. Internet Printing	25
5.1. Trust Relationships	26
5.2. Use of Attribute Certificates	27
5.3. IPP and the Authorization Descriptive Model	28
6. Electronic Commerce	29
6.1. Model Description	30
6.1.1. Identification of Components	30
6.1.2. Identification of Contractual Relationships	31
6.1.3. Identification of Trust Relationships	32
6.1.3.1. Static Trust Relationships	33
6.1.3.2. Dynamic Trust Relationships	35

6.1.4. Communication Model	35
6.2. Multi Domain Model	37
6.3. Requirements	38
7. Computer Based Education and Distance Learning	40
7.1. Model Description	40
7.1.1. Identification of Components	40
7.1.2. Identification of Contractual Relationships	41
7.1.3. Identification of Trust Relationships	43
7.1.4. Sequence of Requests	44
7.2. Requirements	46
8. Security Considerations	47
Glossary	47
References	48
Authors' Addresses	50
Full Copyright Statement	53

1. Introduction

This document is one of a series of three documents under consideration by the AAAarch RG dealing with the authorization requirements for AAA protocols. The three documents are:

- AAA Authorization Framework [2]
- AAA Authorization Requirements [3]
- AAA Authorization Application Examples (this document)

In this memo, we examine several important Internet applications that require authorization. For each application, we present a model showing how it might do authorization and then map that model back to the framework presented in [2]. We then present the authorization requirements of the application as well as we presently understand them. The requirements presented in this memo have been collected together, generalized, and presented in [3].

The intent of this memo is to validate and illustrate the framework presented in [2] and to motivate the requirements presented in [3]. This work is intended to be in alignment with the work of the various working groups responsible for the authorization applications illustrated. This memo should not, however, be regarded as authoritative for any of the applications illustrated. Where authoritative documents exist or are in development, they are listed in the references at the end of this document.

The work for this memo was done by a group that originally was the Authorization subgroup of the AAA Working Group of the IETF. When the charter of the AAA working group was changed to focus on MobileIP and NAS requirements, the AAAarch Research Group was chartered within the IRTF to continue and expand the architectural work started by the Authorization subgroup. This memo is one of four which were created by the subgroup. This memo is a starting point for further work within the AAAarch Research Group. It is still a work in progress and is published so that the work will be available for the AAAarch subgroup and others working in this area, not as a definitive description of architecture or requirements.

This document uses the terms 'MUST', 'SHOULD' and 'MAY', and their negatives, in the way described in RFC 2119 [4].

2. PPP Dialin with Roaming

In this section, we present an authorization model for dialin network access in terms of the framework presented in [2]. Included in the model are the multi-domain considerations required for roaming [5]. Detailed requirements for network access protocols are presented in [6].

2.1. Descriptive Model

The PPP dialin application uses the pull sequence as discussed in [2]. The roaming case uses the roaming pull sequence, also discussed in [2]. This sequence is redrawn using dialin roaming terminology in figure 1, below.

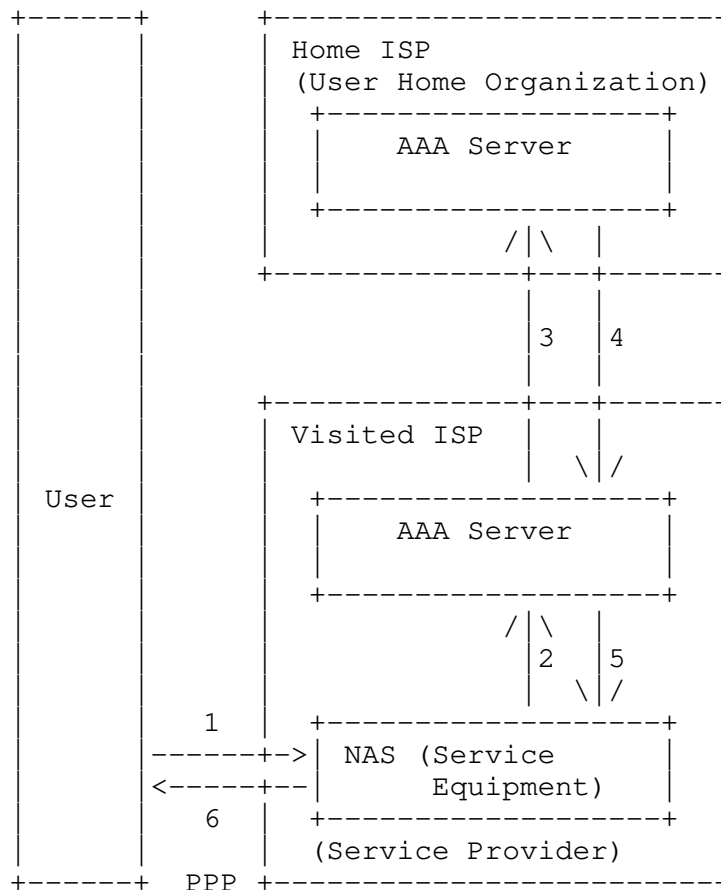


Fig. 1 -- Dialin Authorization
Based on Roaming Pull Sequence

In this model, the User dials in to a Network Access Server (NAS) provided by the visited (or foreign) ISP (the Service Provider in the general model). The User is authenticated using a protocol such as PAP, CHAP, or EAP which is encapsulated in PPP frames (1). Because the User has not yet gained access to the network, he or she cannot send IP datagrams to a AAA server. At this point, the User can only communicate with the NAS (Service Equipment). The NAS forwards the User's authentication/ authorization request including the Network Access Identifier (NAI) [7] to a AAA server in its own domain via RADIUS [8] or a successor AAA protocol (2). The visited ISP's AAA server examines the realm from the NAI and forwards the request to the User's home domain AAA server (3). The home domain AAA server authenticates the user and authorizes access according to a roaming agreement. The home domain AAA server may return service parameters

(e.g. Idle-Timeout) to the visited ISP's AAA server (4) which forwards them to the NAS, possibly adding additional service parameters (5). The NAS completes PPP session initialization (6).

In the future, this model may be expanded in several ways [9]. For instance, Authentication and Authorization may be done in separate passes using different servers in order to support specialized forms of authentication. Or to better support roaming, a broker may be inserted between the visited ISP and the home ISP. Or authorization may be supported based on other identifiers such as the caller ID and called ID obtained from the PSTN (e.g., using ANI and DNIS).

2.2. Authorization Requirements

The following requirements are identified in [9] for authorizing PPP dialin service using roaming.

- Authorization separate from authentication should be allowed when necessary, but the AAA protocol MUST allow for a single message to request both authentication and authorization.
- The AAA protocol MUST be "proxyable", meaning that a AAA Server or PDP MUST be able to forward the request to another AAA Server or PDP, which may or may not be within the same administrative domain.
- The AAA protocol MUST allow for intermediate brokers to add their own local Authorization information to a request or response.
- When a broker is involved, the protocol MUST provide end to end security.
- The broker MUST be able to return a forwarding address to a requester, allowing two nodes to communicate together.
- The protocol MUST provide the following features (per user session):
 1. One Authentication, One Authorization
 2. One Authentication, Multiple Authorization
 3. Multiple Authentication, Multiple Authorization

3. Mobile-IP

The Mobile-IP protocol is used to manage mobility of an IP host across IP subnets [10]. Recent activity within the Mobile-IP Working Group has defined the interaction between Mobile-IP and AAA in order to provide:

- Better scaling of security associations
- Mobility across administrative domain boundaries
- Dynamic assignment of Home Agent

The Mobile IP protocol, as defined in [10], works well when all mobile nodes belong to the same administrative domain. Some of the current work within the Mobile IP Working Group is to allow Mobile IP to scale across administrative domains. This changes the trust model that is currently defined in [10].

The requirements for Mobile-IP authorization are documented in [11]. In this section, we develop a multi-domain model for Mobile-IP authorization and present it in the terms of the framework presented in [2].

Figure 2 depicts the new AAA trust model for Mobile-IP. In this model each network contains mobile nodes (MN) and a AAA server (AAA). Each mobility device shares a security association (SA) with the AAA server within its own home network. This means that none of the mobility devices initially share a security association. Both administrative domains' AAA servers can either share a security association, or can have a security association with an intermediate broker.




Fig. 2 -- Mobile-IP AAA Trust Model

Figure 3 provides an example of a Mobile-IP network that includes AAA. In the integrated Mobile-IP/AAA Network, it is assumed that each mobility agent shares a security association between itself and its local AAA server. Further, the Home and Foreign AAA servers both share a security association with the broker's AAA server. Lastly, it is assumed that each mobile node shares a trust relationship with its home AAA Server.

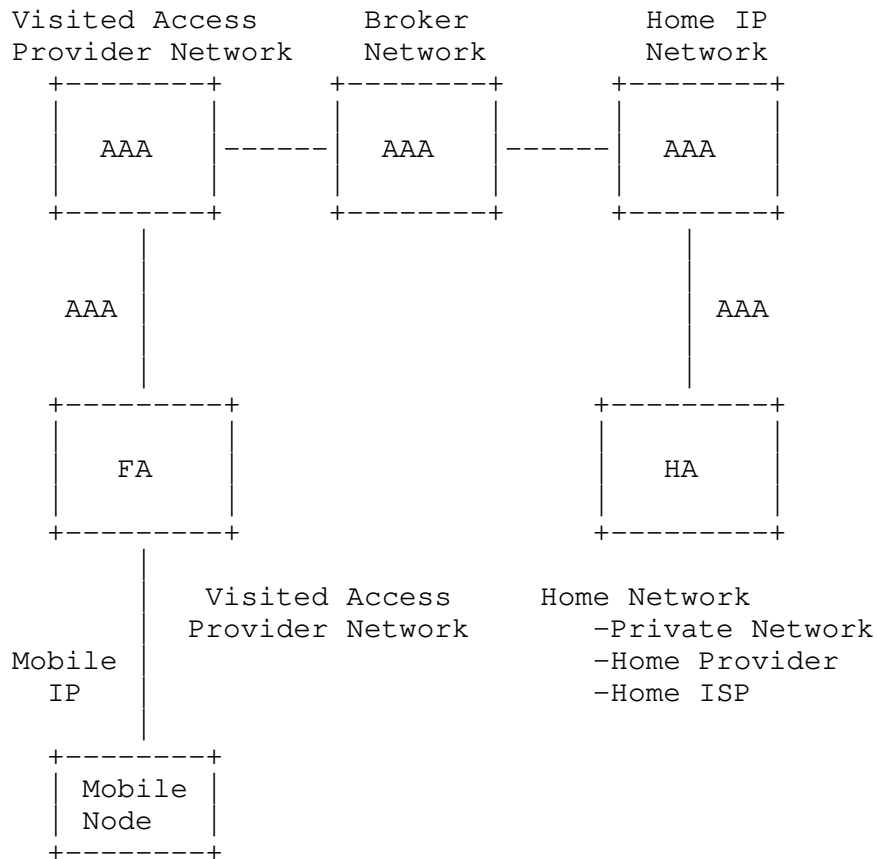


Fig. 3 -- General Wireless IP Architecture for Mobile-IP AAA

In this example, a Mobile Node appears within a foreign network and issues a registration to the Foreign Agent. Since the Foreign Agent does not share any security association with the Home Agent, it sends a AAA request to its local AAA server, which includes the authentication information and the Mobile-IP registration request. The Mobile Node cannot communicate directly with the home AAA Server for two reasons:

- It does not have access to the network. The registration request is sent by the Mobile Node to request access to the network.
- The Mobile Node may not have an IP address, and may be requesting that one be assigned to it by its home provider.

The Foreign AAA Server will determine whether the request can be satisfied locally through the use of the Network Access Identifier [7] provided by the Mobile Node. The NAI has the format of user@realm and the AAA Server uses the realm portion of the NAI to identify the Mobile Node's home AAA Server. If the Foreign AAA Server does not share any security association with the Mobile Node's home AAA Server, it may forward the request to its broker. If the broker has a relationship with the home network, it can forward the request, otherwise a failed response is sent back to the Foreign AAA Server.

When the home AAA Server receives the AAA Request, it authenticates the user and begins the authorization phase. The authorization phase includes the generation of:

- Dynamic Session Keys to be distributed among all Mobility Agents
- Optional Dynamic assignment of a Home Agent
- Optional Dynamic assignment of a Home Address (note this could be done by the Home Agent).
- Optional Assignment of QOS parameters for the Mobile Node [12]

Once authorization is complete, the home AAA Server issues an unsolicited AAA request to the Home Agent, which includes the information in the original AAA request as well as the authorization information generated by the home AAA server. The Home Agent retrieves the Registration Request from the AAA request and processes it, then generates a Registration Reply that is sent back to the home AAA server in a AAA response. The message is forwarded through the broker back to the Foreign AAA server, and finally to the Foreign Agent.

The AAA servers maintain session state information based on the authorization information. If a Mobile Node moves to another Foreign Agent within the foreign domain, a request to the foreign AAA server can immediately be done in order to immediately return the keys that were issued to the previous Foreign Agent. This minimizes an additional round trip through the internet when micro mobility is involved, and enables smooth hand-off.

3.1. Relationship to the Framework

Mobile-IP uses the roaming pull model described in [2]. The Mobile Node is the User. The Foreign Network is the Service Provider with the Foreign Agent as the Service Equipment. The Home Network is the User Home Organization. Note that the User Home Organization operates not only a AAA Server, but also the Home Agent. Note, also, that a broker has been inserted between the Service Provider and the User Home Organization.

3.2. Minimized Internet Traversal

Although it would have been possible for the AAA interactions to be performed for basic authentication and authorization, and the Registration flow to be sent directly to the Home Agent from the Foreign Agent, one of the key Mobile-IP AAA requirements is to minimize Internet Traversals. Including the Registration Request and Replies in the AAA messages allows for a single traversal to authenticate the user, perform authorization and process the Registration Request. This streamlined approach is required in order to minimize the latency involved in getting wireless (cellular) devices access to the network. New registrations should not increase the connect time more than what the current cellular networks provide.

3.3. Key Distribution

In order to allow the scaling of wireless data access across administrative domains, it is necessary to minimize the security associations required. This means that each Foreign Agent does not share a security association with each Home Agent on the Internet. The Mobility Agents share a security association with their local AAA server, which in turn shares a security association with other AAA servers. Again, the use of brokers, as defined by the Roaming Operations (roamops) Working Group, allows such services to scale by allowing the number of relationships established by the providers to be reduced.

After a Mobile Node is authenticated, the authorization phase includes the generation of Sessions Keys. Specifically, three keys are generated:

- k1 - Key to be shared between the Mobile Node and the Home Agent
- k2 - Key to be shared between the Mobile Node and the Foreign Agent
- k3 - Key to be shared between the Foreign Agent and the Home Agent

Each Key is propagated to each mobility device through the AAA protocol (for the Foreign and Home Agent) and via Mobile-IP for the Mobile Node (since the Mobile Node does not interface directly with the AAA servers).

Figure 4 depicts the new security associations used for Mobile-IP message integrity using the keys derived by the AAA server.




Fig. 4 -- Security Association after Key Distribution

Once the session keys have been established and propagated, the mobility devices can exchange registration information directly without the need of the AAA infrastructure. However the session keys have a lifetime, after which the AAA infrastructure must be used in order to acquire new session keys.

3.4. Mobile-IP Authorization Requirements

To summarize, Mobile-IP has the following authorization requirements:

1. Mobile-IP requires an AAA protocol that makes use of the pull model.
2. Mobile-IP requires broker support, and data objects must contain data integrity and confidentiality end-to-end. This means that neither the broker nor any other intermediate AAA node should be able to decrypt the data objects, but they must be able to verify the objects' validity.
3. Authorization includes Resource Management. This allows the AAA servers to maintain a snapshot of a mobile node's current location, keying information, etc.

4. Due to the nature of the service being offered, it is imperative that the AAA transaction add minimal latency to the connect time. Ideally, the AAA protocol should allow for a single round trip for authentication and authorization.
5. If the AAA protocol allows for the Mobile-IP registration messages to be embedded within the authentication/authorization request, this will further reduce the number of round trips required and hence reduce the connect time.
6. It must be possible to pass Mobile-IP specific key management data along with the authorization data. This allows the AAA server to act as a Key Distribution Center (KDC).
7. It must be possible to pass other application-specific data units such as home agent selection and home address assignment to be carried along with the authorization data units.
8. The authorization response should allow for diffserv (QOS) profiles, which can be used by the mobility agents to provide some quality of service to the mobile node.
9. The AAA protocol must allow for unsolicited messages to be sent to a "client", such as the AAA client running on the Home Agent.

4. Bandwidth Broker

This section describes authorization aspects derived from the Bandwidth Broker architecture as discussed within the Internet2 Qbone BB Advisory Council. We use authorization model concepts to identify contract relationships and trust relationships, and we present possible message exchanges. We will derive a set of authorization requirements for Bandwidth Brokers from our architectural model. The Internet 2 Qbone BB Advisory Council researches a single and multi-domain implementation based on 2-tier authorization concepts. A 3-tier model is considered as a future work item and therefore not part of this description. Information concerning the Internet 2 Bandwidth Broker work and its concepts can be found at:

<http://www.merit.edu/working.groups/i2-qbone-bb>

The material in this section is based on [13] which is a work in progress of the Internet2 Qbone BB Advisory Council.

4.1. Model Description

The establishment of a model involves four steps:

1. identification of the components that are involved and what they are called in this specific environment,
2. identification of the relationships between the involved parties that are based on some form of agreement,
3. identification of the relationships that are based on trust, and
4. consideration of the sequence of messages exchanged between components.

4.2. Components of the Two-Tier Model for Bandwidth Brokerage

We will consider the components of a bandwidth broker transaction in the context of the conceptual entities defined in [2]. The bandwidth broker two-tier model recognizes a User and the Service Provider controlling the Service Equipment.

The components are as follows:

- The Service User (User) -- A person or process willing to use certain level of QoS by requesting the allocation of a quantifiable amount of resource between a selected destination and itself. In bandwidth broker terms, the User is called a Service User, capable of generating a Resource Allocation Request (RAR).
- The Bandwidth Broker (Service Provider) -- a function that authorizes allocation of a specified amount of bandwidth resource between an identified source and destination based on a set of policies. In this context we refer to this function as the Bandwidth Broker. A Bandwidth Broker is capable of managing the resource availability within a network domain it controls.

Note: a 3-tier model involving a User Home Organization is recognized in [13], however its development is left for future study and therefore it is not discussed in this document.

4.3. Identification of Contractual Relationships

Authorizations to obtain bandwidth are based on contractual relationships. In both the single and multi-domain cases, the current Bandwidth Broker model assumes that a User always has a contractual relationship with the service domain to which it is connected.

4.3.1. Single-Domain Case

In the single-domain case, the User has a contract with a single Service Provider in a single service domain.




Fig. 5 -- Two-Tier Single Domain Contractual Relationships

4.3.2. Multi-Domain Case

In the multi-domain case, the User has a contract with a single Service Provider. This Service Provider has a contract with neighboring Service Providers. This model is used when independent autonomous networks establish contracts with each other.

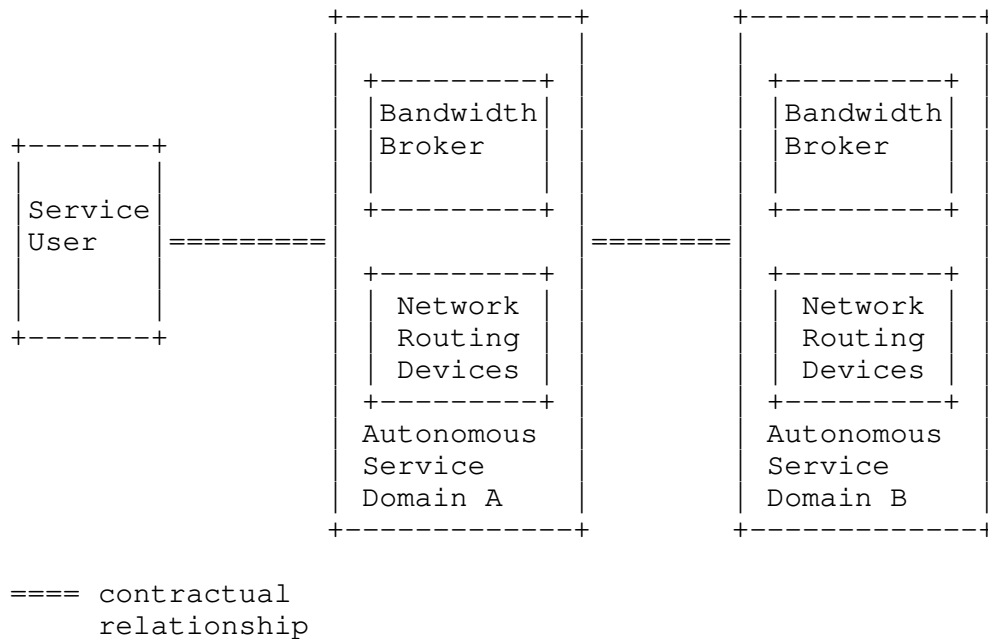


Fig. 6 -- Two-Tier Multi-Domain Contractual Relationships

4.4. Identification of Trust Relationships

Contractual relationships may be independent of how trust, which is necessary to facilitate authenticated and possibly secure communication, is implemented. There are several alternatives in the Bandwidth Broker environment to create trusted relationships. Figures 7 and 8 show two alternatives that are options in the two-tier Bandwidth Broker model.

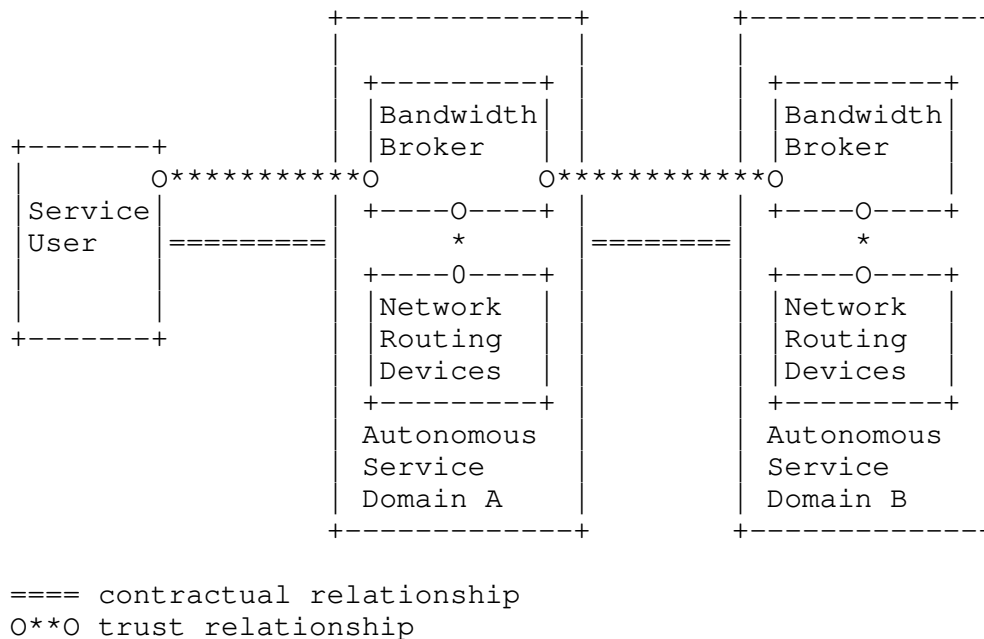


Fig. 7 -- Two-Tier Multi-Domain Trust Relationships, alt 1

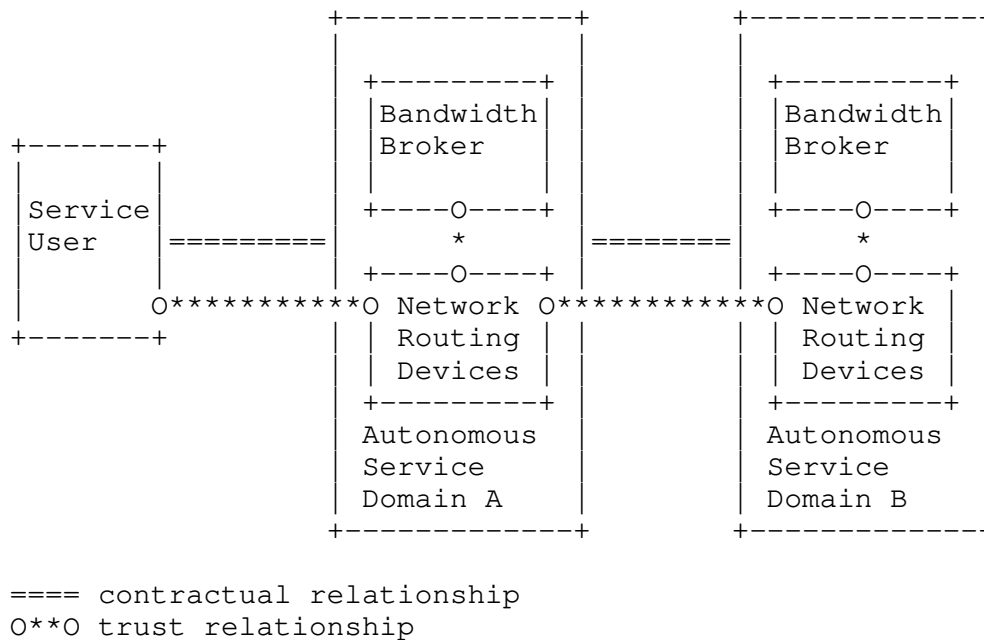
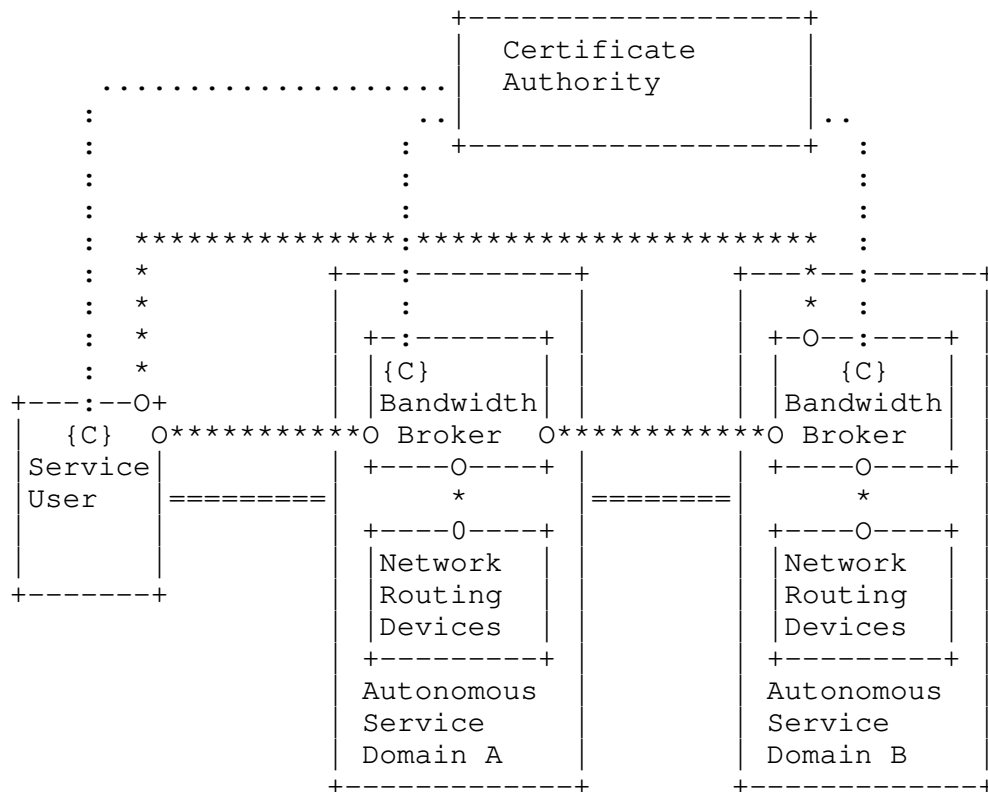


Fig. 8 -- Two-Tier Multi-Domain Trust Relationships, alt 2

Although [13] does not recommend specifics regarding this question, the document recognizes the need for trust relationships. In the first model, a trust relationship, based on some form of authentication method, is created between the User and the Bandwidth Broker and among Bandwidth Brokers. In the second model, which enjoys some popularity in enterprise networks, the trust relationship may be established via the wiring closet and the knowledge of which physical router port or MAC address is connected to which user. The router-Bandwidth Broker relationship may be established physically or by some other authentication method or secure channel.

A Certificate Authority (CA) based trust relationship is shown in figure 9. In this figure, a CA signs public key certificates, which then can be used in encrypted message exchanges using public keys that are trusted by all involved. As a first step, each involved party must register with the CA so it can join a trust domain. The Router-Bandwidth Broker relationship may be established as described in the two previous figures. An interesting observation regarding this kind of model is that the bandwidth broker in domain B may route information to the user via the bandwidth broker in domain A without BB1 being able to read the information (using end-to-end security). This model creates a meshed trust relationship via a tree like CA structure.



```
==== contractual relationship
0**0 trust relationship
{C}. certification process
```

Fig. 9 -- Two-Tier Multi-Domain Trust Relationships, alt 3

4.5. Communication Models and Trust Relationships

When describing the Bandwidth Broker communication model, it is important to recognize that trust relationships between components must ensure secure and authenticated communication between the involved components. As the Internet 2 Qbone Bandwidth Broker work does not recommend any particular trust relationship model, we make the same assumptions as [13]. In theory, the trust model and communication model can be independent, however communication efficiency will determine the most logical approach.

4.6. Bandwidth Broker Communication Models

4.6.1. Concepts

The current Internet 2 Qbone Bandwidth Broker discussion describes a two-tier model, where a Bandwidth Broker accepts Resource Allocation Requests (RAR's) from users belonging to its domain or RAR's generated by upstream Bandwidth Brokers from adjacent domains. Each Bandwidth Broker will manage one service domain and subsequently provide authorization based on a policy that decides whether a request can be honored.

4.6.1.1. Intra-Domain Authorization

Admission Authorization or Connection Admission Control (CAC) for intra-domain communication is performed using whatever method is appropriate for determining availability of resources within the domain. Generally a Bandwidth Broker configures its service domain to certain levels of service. RAR's are subsequently accommodated using a policy-based decision.

4.6.1.2. Inter-Domain Authorization

Service Level Specifications (SLS's) provide the basis for handling inter-domain bandwidth authorization requests. A Bandwidth Broker monitors both the state of its network components and the state of its connections to neighboring networks. SLS's are translations of SLA's established between Autonomous Service Domains. Each Bandwidth Broker will initialize itself so it is aware of existing SLS's. SLS's are established in a unidirectional sense. Two SLS's must govern a bi-directional connection. SLS's are established on the level of aggregate data-flows and the resources (bandwidth) provisioned for these flows.

A Bandwidth Broker may honor an inter-domain RAR by applying policy decisions determining that a particular RAR does fit into a pre-established SLS. If successful, the Bandwidth Broker will authorize the usage of the bandwidth. If unsuccessful, the Bandwidth Broker may deny the request or approve the request after it has re-negotiated the SLS with its downstream Bandwidth Broker.

A separate Policy Manager may be involved in the CAC decision. The Internet 2 Qbone Bandwidth Broker discussion recognizes an ideal environment where Bandwidth Brokers and Policy Managers work together to provide CAC using integrated policy services [13].

4.6.2. Bandwidth Broker Work Phases

The Internet 2 Qbone Bandwidth Broker discussion proposes development of the Bandwidth Broker model in several phases:

- Phase 0: Local Admission. RAR's are only handled within a local domain. SLS's are pre-established using manual methods (fax, e-mail).
- Phase 1: Informed Admission. RAR's spanning multiple domains are authorized based on information obtained from one or more Bandwidth Brokers along the path.
- Phase 2: Dynamic SLS admission. Bandwidth Brokers can dynamically set up new SLS's.

Although the local admission case is addressed, the current Internet 2 Qbone Bandwidth Broker work is currently concerned with solving multi-domain problems in order to allow individual Bandwidth Brokers to inter-operate as identified in phase 0 or 1.

4.6.3. Inter-Domain Signaling


4.6.3.1. Phase 0

In phase 0 implementations, no electronic signaling between Bandwidth Brokers is performed and SLS negotiation will be performed manually (phone, email etc) by network operators. An RAR is only handled within the domain and may originate from a User or ingress router.


4.6.3.2. Phase 1

Here a CAC decision is made on information obtained from downstream Bandwidth Brokers. This information could come from the next hop Bandwidth Broker or all Bandwidth Brokers downstream to the destination.

Two fundamental signaling approaches between Bandwidth Brokers have been identified for the Informed Admission case. These are illustrated in figure 10.



A) End-to-end signaling



B) Immediate response signaling.

Fig. 10 -- Fundamental Signalling Approaches

- End to End signaling. An RAR from a User to BB1 is forwarded to BB2 (1). BB2 will forward the request to BB3 (2). If BB3 is the destination of the request, BB3 will authorize the request and reply to BB2 (3). BB2 will then reply to BB1 (4), and BB1 will send a Resource Allocation Answer (RAA) back to the User to complete the authorization.
- Immediate response signaling. This is the case where BB1 will want to authorize an RAR from its domain and forwards the authorization request to BB2 (1). If BB2 approves, the response is immediately returned to BB1 (2). BB1 will send an RAA back to the User. If the authorization was positive BB2 will forward subsequently a request to the next BB, BB3 (3). BB3 authorizes the request and responds to BB2 (4). If the response is negative (5), BB2 will cancel the authorization it previously issued to BB1 (6) and this will result in a cancellation from BB1 to the user (7). In this case the RAA authorization is valid until revoked by 7.

4.6.4. Bandwidth Broker Communication Architecture

Figure 11 shows components of the discussed Bandwidth Broker architecture with its interfaces.

- An intra-domain interface allows communication with all the service components within the network that the Bandwidth Broker controls.
- An inter-domain interface allows communication between Bandwidth Brokers of different autonomous networks.
- A user/application interface allows the Bandwidth Broker to be managed manually. Requests can be sent from the User or a host application.
- A policy manager interface allows implementation of complex policy management or admission control.
- A routing table interface allows the Bandwidth Broker to understand the network topology.
- An NMS interface allows coordination of network provisioning and monitoring.

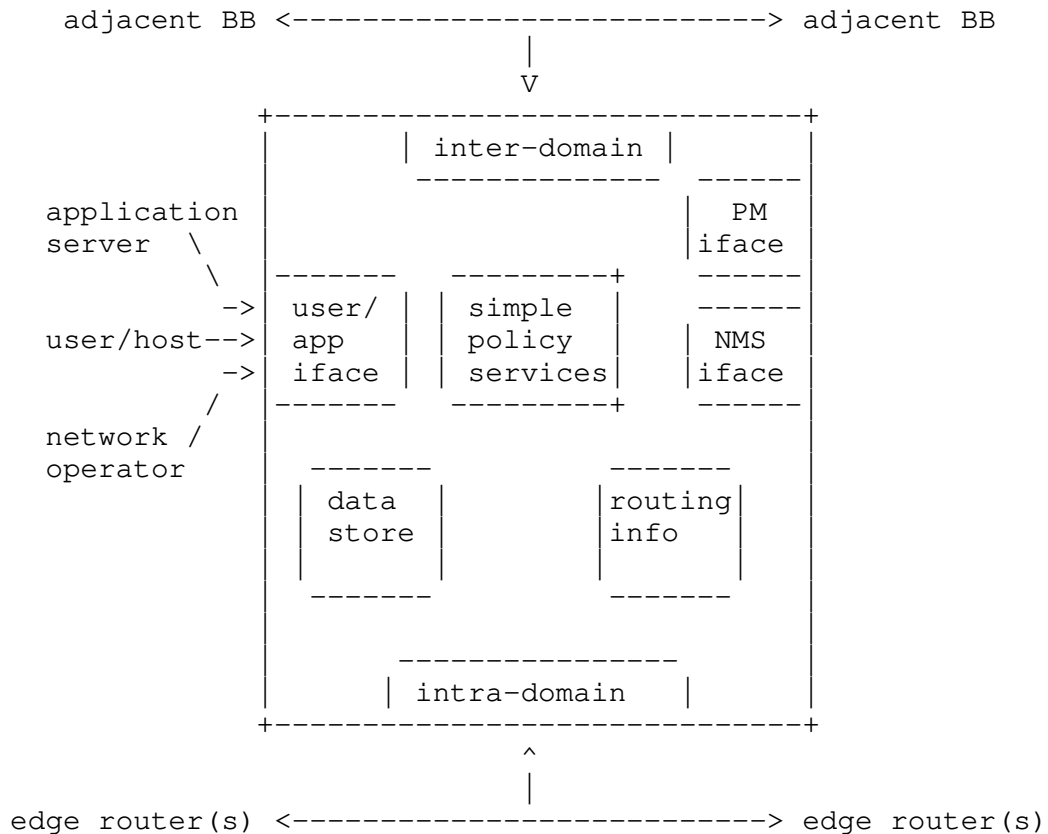


Fig. 11 -- Bandwidth Broker Architecture

4.6.5. Two-Tier Inter-Domain Bandwidth Broker Communication Model

4.6.5.1. Session Initialization

Before Bandwidth Brokers can configure services between two adjacent domains, they have to establish and initialize a relationship. No authentication is used; therefore any trust relationship is implicit. Part of the initialization is an exchange of topology information (list of adjacent Bandwidth Brokers).

4.6.5.2. Service Setup

The Bandwidth Broker must first be configured in regard to agreed bi-lateral service levels. All resources allocated to a particular level of provisioned service must be reserved in each domain.

A Service Setup Request (SSR) is generated (on demand by the operator or at startup of the system) and forwarded to a downstream Bandwidth Broker. The downstream Bandwidth Broker will check the

consistency with its own service level specifications and respond with Setup Answer message (SA) agreements. This message exchange confirms and identifies pre-established service authorization levels.

4.6.5.3. Service Cancellation

A Service Cancellation (SC) message may cancel a service authorization. This message may be initiated by the operator or by an expiration date. A Cancellation Answer (CA) is returned.

4.6.5.4. Service Renegotiation

An (optional) Service-Renegotiation message (SR) may allow a Bandwidth Broker to re-negotiate an existing service. This message may be initiated by the operator or automatically when a certain threshold is reached. Renegotiations happen within the margins of a pre-established authorization.

4.6.5.5. Resource Allocation Request and Resource Allocation Answer

An RAR allocates a requested level of service on behalf of the User and when available it will decide on the admittance of a certain User to the service. A Bandwidth Broker may receive an RAR via either the intra-domain or inter-domain interface. The RAR must refer to the Service SetUp Identification (SSU_ID), which binds a request to a certain authorization. A Resource Allocation Answer (RAA) confirms or rejects a request or it may indicate an "in progress" state.

4.6.5.6. Session Maintenance

A certain level of session maintenance is required to keep Bandwidth Brokers aware of each other. This must be implemented using time-outs and keep-alive messages. This will help Bandwidth Brokers to notice when other Bandwidth Brokers disappear.

4.6.5.7. Intra-domain Interface Protocol

The Intra-domain interface protocol used between a Bandwidth Broker and the routers it controls may be COPS, SNMP, or Telnet Command Line Interface.

4.7. Requirements

From the above descriptions we derive the following requirements.

- The Authorization mechanism may require trust relationships to be established before any requests can be made from the User to the Service Provider. Currently trust relationship establishment is implicit.
- A confirmation of authorization is required in order to initialize the system.
- A negation of static authorization is required to shut down certain services.
- A renegotiation of static authorization is required to alter services (SLS's).
- Dynamic authorization requests (RAR) must fit into pre-established static authorizations (SLS's).
- Dynamic authorization requests (RAR) may be answered by an "in progress state" answer.
- Provisions must be made to allow reconstruction of authorization states after a Bandwidth Broker re-initializes.

5. Internet Printing

The Internet Printing Protocol, IPP [14], has some potentially complex authorization requirements, in particular with the "print-by-reference" model. The following attempts to describe some possible ways in which an authorization solution for this aspect of IPP might work, and to relate these to the framework described in [2]. This is not a product of the IPP working group, and is meant only to illustrate some issues in authorization in order to establish requirements for a "generic" protocol to support AAA functions across many applications.

IPP print-by-reference allows a user to request a print service to print a particular file. The user creates a request to print a particular file on a printer (or one of a group of printers). The key aspect is that the request includes only the file name and not the file content. The print service must then read the file from a file server prior to printing. Both the file server and the print server must authorize the request. Once initiated, printing will be done without intervention of the user; i.e., the file will be sent directly to the print service rather than through the user to the printer.

5.1. Trust Relationships

The assumption is that the Printer and File Server may be owned and operated by different organizations. There appear to be two models for how "agreements" can be set up.

1. User has agreement with Print Server; Print Server has agreement with File Server.
2. User has agreements with both File and Print Server directly.

In case 1, the user has a trust relationship with the Print Service AAA Server. The Printer forwards the request to the File Server. The File Server authorizes the Printer and determines if the Printer is allowed access to the file. Note that while there may be some cases where a Print Server may on its own be allowed access to files (perhaps some "public files", or that can only be printed on certain "secure" printers), it is normally the case that files are associated with users and not with printers. This is not a good "generic" model as it tends to make the print service an attractive point of attack.




Fig. 12 -- Case 1
 User authorizes with Print Service.
 Printer authorizes with File Service.

In case 2, the user must have a trust relationship with both the file and print services so that each can verify the service appropriate to the User. In this case, the User first contacts the File Service AAA Server and requests that it enable authorization for the Print

Service to access the file. This might be done in various ways, for example the File Service AAA Server may return a token to the User which can (via the Print Service) be presented to the File Server to enable access.




Fig. 13 -- Case 2

User authorizes File and Print Service.
Must create binding for session between
Print Service and File Service.

5.2. Use of Attribute Certificates in Print-by-Reference

The print-by-reference case provides a good example of the use of attribute certificates as discussed in [2]. If we describe case 2 above in terms of attribute certificates (ACs) we get the diagram shown in figure 14.




Fig. 14 -- Using Attribute Certificates in IPP Authorization

In this case, the User gets an AC from the File Service's AAA Server which is signed by the File Service AAA Server and contains a set of attributes describing what the holder of the AC is allowed to do. The User then authorizes with the Print Service AAA Server and passes the AC in the authorization request. The Printer establishes a session with the File Server, passing it the AC. The File Server trusts the AC because it is signed by the File Service AAA Server and allows (or disallows) the session.

It is interesting to note that an AC could also be created and signed by the User, and passed from the Print Server to the File Server. The File Server would need to be able to recognize the User's signature. Yet another possibility is that the Print Service AAA Server could simply authenticate the User and then request an AC from the File Service AAA Server.

5.3. IPP and the Authorization Descriptive Model

The descriptive model presented in [2] includes four basic elements: User, User Home Organization, Service Provider AAA Server, and Service Equipment.

Mapping these to IPP, the User is the same, the User Home Organization (if included) is the same. The Service Provider AAA Server and the Service Equipment are expected to be closely coupled on the same processor. In other words, the interface between the

Print Service AAA Server and the Printer as well as that between the File Service AAA Server and the File Server is an internal one that will not require a formal protocol (although some standard API might be useful).

The concept of a Resource Manager (see [2]) has some interesting twists relative to IPP. Once started, the user is not involved in the service, but until printing is complete it seems useful that any of the parties in the authorization process be allowed to query for status or to cancel the print session. The user needs a way to "bind" to a particular session, and may have to reauthorize to be allowed to access Resource Manager information.

6. Electronic Commerce

This section describes the authorization aspects of an e-commerce architecture typically used in Europe. We will use this model to identify contractual and trust relationships and message exchanges. We will then identify a set of authorization requirements for e-commerce.

Whereas most e-commerce protocols focus on authentication and message integrity, e-commerce exchanges as described by the Internet Open Trading Protocol (trade) Working Group in [15] also involve authorization. This section will examine one e-commerce protocol called SET (Secure Electronic Transaction) that provides for credit and debit card payments. We will analyze the authorization aspects from an architectural viewpoint. We will apply concepts and terms defined in [2].

We are not here proposing SET as a standard authorization protocol. Rather, we are examining the SET model as a way of understanding the e-commerce problem domain so that we can derive requirements that an authorization protocol would have to meet in order to be used in that domain.

E-commerce protocols and mechanisms such as those described in [16] may not only be important to allow customers to shop safely in Cyberspace, but may also be important for purchases of Internet services as well. With emerging technologies allowing Internet transport services to be differentiated, an inherently more complex pricing model will be required as well as additional payment methods. Flexible authorization of services will be an important aspect to allow, for example, globally roaming users ad hoc allocation of premium bandwidth with an ISP who is authorized to accept certain credit card brands.

6.1. Model Description

The establishment of a model involves four steps:

1. identification of the components that are involved and what they are called in this specific environment,
2. identification of the relationships between the involved parties that are based on some form of agreement,
3. identification of the relationships that are based on trust, and
4. consideration of the sequence of messages exchanged between components.

6.1.1. Identification of Components

We will consider the components of an electronic commerce transaction in the context of the conceptual entities defined in [2].

- The Cardholder (User) -- the person or organization that is to receive and pay for the goods or services after a request to purchase has been received. In SET terms this is called a Cardholder.
- The Issuer (User Home Organization) -- the financial organization that guarantees to pay for authorized transactions to purchase goods or services on behalf of the User when using a debit or credit card it issues. The financial organization (typically a bank or Brand Organization) will transfer money from the user account to the account the party to which the User instructs it to send the payment. The issued card authorizes the User to use the card for payments to merchants who are authorized to accept the card. In SET terms this organization is called the Issuer. This organization is considered "home" to the Cardholder.
- The Merchant (Service Provider) -- the organization from whom the purchase is being made and who is legally responsible for providing the goods or services and receives the benefit of the payment made. In SET terms this organization is called a Merchant. The Cardholder is considered to be "foreign" to the Merchant.
- The Acquirer (Broker) -- the organization that processes credit or debit card transactions. Although in reality this function may be rather complex and may span several organizations, we will simply assume this organization to be a Brand Organization fulfilling the role of the Acquirer as defined in SET. The Acquirer establishes an account with the Merchant. The Acquirer operates a Payment Gateway that will accept payment authorization requests from

authorized merchants and provide responses from the issuer. The Acquirer will forward an authorization request to the Issuer. The Acquirer is considered "home" to the Merchant.

As the SET document [16] notes, a Brand Organization (credit card organization) may handle both the Issuer function and Acquirer function that operates a Payment Gateway. For simplicity, we therefore assume that the authorization role of Broker (Acquirer) and User Home Organization (Issuer) both belong to the Brand Organization.

In order to be more descriptive we now use the SET terms. In the requirements section these terms are mapped back into the authorization framework terms again.

6.1.2. Identification of Contractual Relationships

Contractual relationships are illustrated in figure 15, below.

- The Cardholder has a contractual relationship with the card Issuer. The Cardholder holds an account with the Issuer and obtains an account number.
- The Merchant has a contractual relationship with the Acquirer. The Merchant obtains a Merchant ID from the Acquirer.
- In the real world there may be no direct contractual relationship between the Issuer and the Acquirer. The contractual relationships allowing an Acquirer to relay a payment authorization request to an Issuer may be very complex and distributed over multiple organizations. For simplicity, however, we assume there are contracts in place allowing an Acquirer to request payment authorization from an Issuer. These contracts are facilitated by the Brand Organization. Therefore, in our simplified example, the Acquirer and Issuer belong to the same Brand Organization. The Acquirer operates a Payment Gateway for which it needs a Bank Identification Number (BIN).

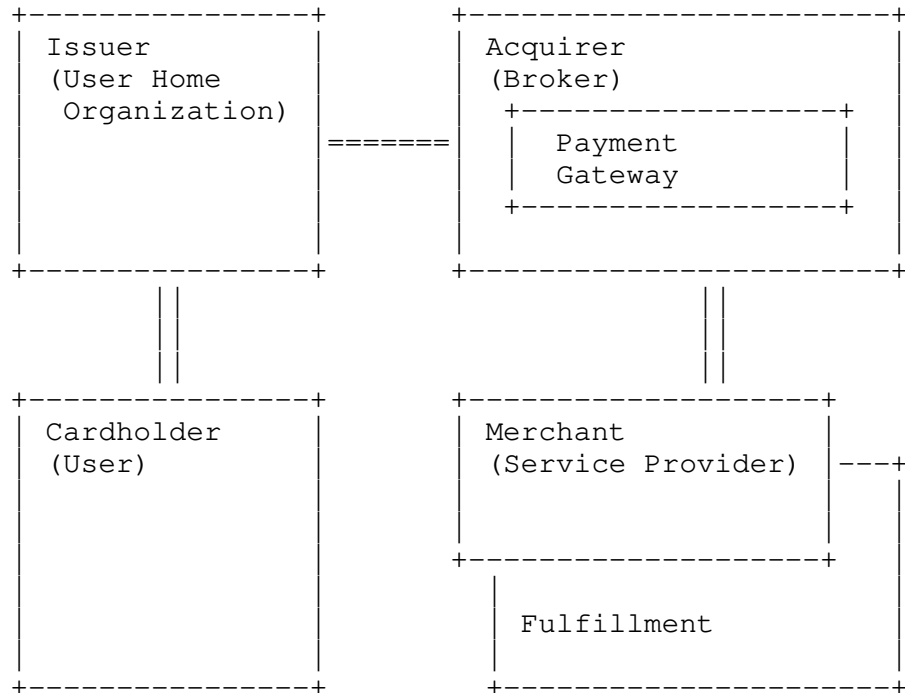


Fig. 15 -- SET Contractual Relationships

6.1.3. Identification of Trust Relationships

It is important to recognize that there are two kinds of trust relationships: static and dynamic trust relationships. Static trust relationships in SET are established by means of a registration process that will request a certificate to be issued to the party that needs to be trusted and authorized to be part of a SET transaction. Dynamic trust is created at the time of a payment transaction and its subsequent authorization request. Note that at the issue phase of a certificate, based on identification and registration, the user of the certificate gets an implicit static authorization and a means of authenticating and securing messages. For this purpose a Certificate Authority (CA) will issue certificates that are used to sign and/or encrypt messages exchanged according to the SET protocol.

6.1.3.1. Static Trust Relationships

In the discussion that follows, refer to figure 16, below.

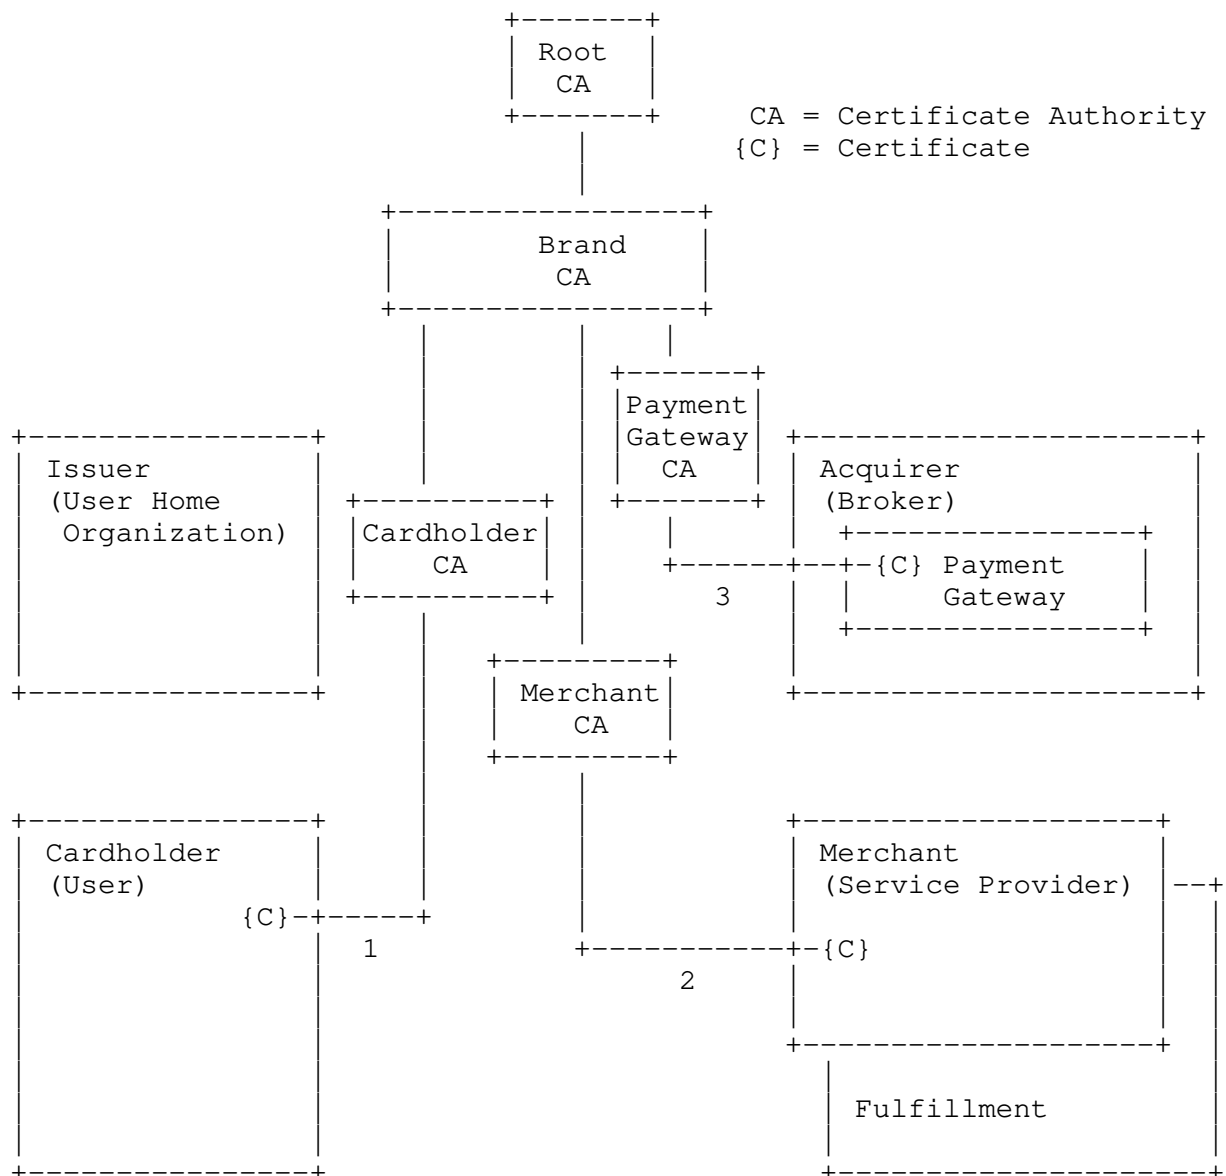


Fig. 16 -- SET Trust Relationships within a Brand Domain

- The Brand Organization operates a Brand CA and is therefore the holder of the common trust within the described domain. All involved parties (Cardholder, Issuer, Merchant and Acquirer) are members of the same trust domain. We will identify three separate

CA's which issue a certificate on behalf of the Issuer, the Acquirer and the Brand Organization. The Brand CA, according to a tree like hierarchy, certifies all underlying CA's. The Brand CA obtains its trust from a single Root Certificate Authority. Before any party can obtain a Certificate from a CA, the party must have some form of contractual relationship.

- After an account has been established with the Issuer, the Cardholder has to register with a Cardholder CA (CCA) through a series of registration steps (1) as defined in the SET protocol. If the CCA approves the registration, the Cardholder will obtain a Cardholder Certificate. The CCA may be operated by the Brand Organization on behalf of the Issuer. The Cardholder Certificate is an electronic representation of the payment card. This process creates a trust relationship between the Cardholder and the Brand. After the cardholder has received the Cardholder Certificate, the Cardholder is authorized to perform payments to an authorized Merchant.
- After the Merchant has obtained a Merchant ID from the Acquirer, the Merchant has to register with the Merchant CA (MCA) through a series of registration steps (2) as defined in the SET protocol. If the MCA approves the registration, the Merchant will obtain a Merchant Certificate. This process creates a trust relationship between the Merchant and the Brand. The MCA may be operated by the Brand Organization on behalf of the Acquirer. After registration, the Merchant is authorized to accept payment requests from Cardholders and to send authorization requests to the Acquirer's Payment Gateway.
- After the Acquirer has obtained a valid Bank Identification Number (BIN), the Acquirer must register with the Payment Gateway CA (PCA) in order to obtain a Payment Gateway Certificate (3). The Payment Gateway Certificate authorizes the Gateway to accept payment authorization requests originating from Merchants within its trust domain.
- The Acquirer and Issuer have a trust relationship via the Brand Organization. The trust relationship is not ensured by procedures or a mechanism defined by SET, as this is a problem solved by agreements between financial organizations facilitating the payment service. Again, for simplicity, we assume that the relationship ensures that payment authorization requests received by the Acquirer's gateway will be forwarded in a secure and efficient way to the Issuer and its response is handled in the same way.

6.1.3.2. Dynamic Trust Relationships

Note that there is no prior established static trust relationship between the Cardholder and the Merchant, as a Cardholder does not have to register with a Merchant or vice versa. The trust relationship is dynamically created during the communication process and is based on the common relationship with the Brand. By means of digital signatures using public key cryptography, the Cardholder's software is able to verify that the Merchant is authorized to accept the Brand Organization's credit card. The merchant is able to verify that the Cardholder has been authorized to use the Brand Organization's credit card.

6.1.4. Communication Model

The purchase request from Cardholder to Merchant and subsequent payment authorization exchange between Merchant and Acquirer is illustrated in figure 17 and described below.

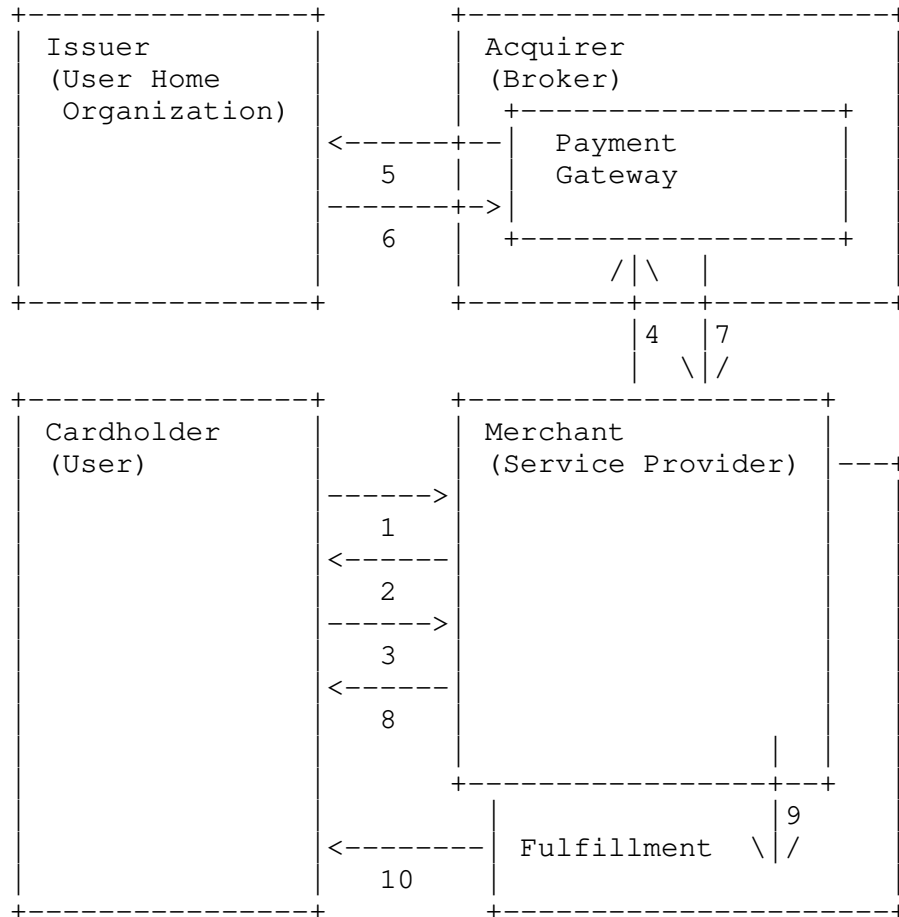


Fig. 17 -- Communication Sequence

1. The Cardholder shops and decides to purchase some goods at merchant.com. The Cardholder has selected a list of goods and the Merchant's software has subsequently prepared an order form for the Cardholder indicating the price, the terms and conditions, and the accepted payment methods. The SET transaction starts at the moment the Cardholder indicates that he or she wants to pay for the goods using a certain payment brand. The Cardholder software sends a request to the Merchant that initiates the payment process.
2. The Merchant checks the order and signs it and returns it to the Cardholder including a certificate from the Acquirer's Gateway that allows the Cardholder to encrypt payment instructions that are only relevant to the Gateway and not to the Merchant (e.g., the Cardholder's credit card information). The Cardholder also includes his or her own certificate.

3. The Cardholder now verifies both certificates (the software has the CA's root certificate). The Cardholder software generates a message containing the order information and the payment instructions that is signed by the Cardholder. Using the Gateway Certificate, it will encrypt the Payment Instruction so that it will only be readable by the Gateway. The Cardholder will include his or her certificate.
4. The Merchant verifies the Cardholder certificate and checks the message integrity. He or she will now process the payment and issue a payment authorization request to the gateway. The payment authorization request contains the Cardholder's certificate and both Merchant certificates.
5. The Gateway verifies the Merchant's signature certificate and that the Merchant signed the authorization request. Next it will obtain the account information and payment instructions and will check the message integrity and the Cardholder's certificate. If everything is in proper order it will send an authorization request to the Issuer via a secure bank network.
6. The issuer returns the authorization.
7. The Acquirer's Gateway generates an authorization response which includes the gateway's certificate.
8. The Merchant checks the authorization response and completes the process by forwarding a purchase response to the Cardholder.
9. The Merchant software authorizes the delivery of the purchased goods.
10. The Cardholder receives the purchased goods.

6.2. Multi Domain Model

In the previous "single" domain case we already assume that there are multiple Cardholders, Merchants, Issuers and Acquirers. However all these parties belong to a single trust domain as there is only a single CCA, MCA and PCA. The trust relationship between multiple cardholders and multiple Issuers go via a single CCA in the same way as the trust relationship between an Acquirer and a Merchant uses the same MCA. The multi-domain case arises when there are multiple domains of CCA's, MCA's and PCA's. In SET these domains reside under a particular Geopolitical CA (GCA) which is illustrated in figure 18.

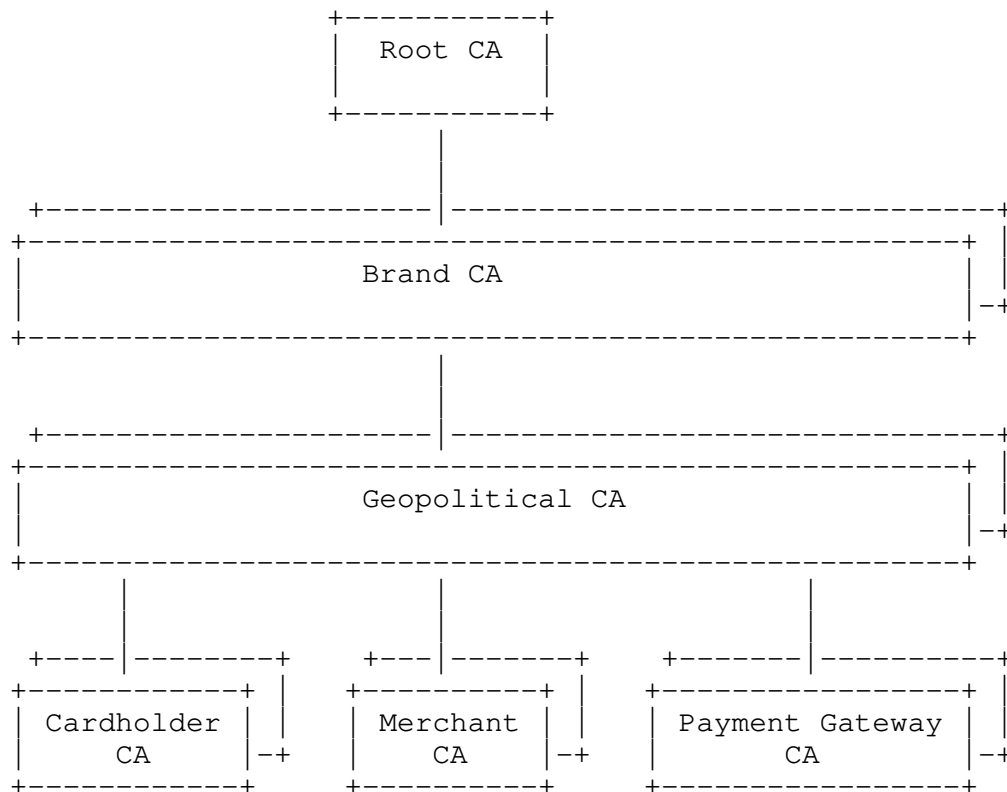


Fig. 18 -- SET Certificate Management Architecture

A GCA may represent a country or region. The architecture defines a trust hierarchy needed to manage and verify SET Certificates as these need to be issued, renewed or revoked. Each geopolitical region may have different policies for issuing, renewing or revoking certificates. However once certificates have been issued, Cardholders and Merchants belonging to different GCA's can still be recognized as belonging to the same Brand. This will allow a European Cardholder to purchase goods in the U.S. The U.S. Acquirer's gateway will recognize that the Cardholder belongs to the same Brand and will therefore accept a payment authorization request.

6.3. Requirements

Many e-commerce environments do not use SET. Other mechanisms exist based on SSL, XML, and S/MIME. Also a mechanism that uses SET only for the payment authorization to the Gateway exists and is known as half SET. However, using the model described in this document, we can derive a fairly comprehensive set of protocol requirements for e-commerce. In these requirements, the SET terms are replaced again by the descriptive model terms:

Cardholder = User
Merchant = Service Provider
Issuer = User Organization
Acquirer = Broker

1. The Authorization mechanism must allow trust relationships to be established before any requests can be made from the User to the Service Provider and from the Service Provider via a Broker to the User Organization. This process will enable the parties to communicate securely by creating an authenticated channel and, by so doing, implicitly authorizing its usage.
2. Upon receipt of any request or response, entities need to be able to verify whether the transmitting party is still authorized to send this request or response.
3. The User must be able to authorize the Service Provider to request an authorization from the User Home Organization.
4. The User must be able to authorize fulfillment of a proposed service offer from the Service Provider.

Other requirements related to the authorization process:

Integrity

5. For any authorization request or response, the receiving party needs to verify that the content of the message has not been altered.

Confidentiality/Privacy

6. The User must be able to pass information relevant to the session authorization process to the User Home Organization via a Broker and the Service Provider without allowing the Broker or the Service Provider to examine its content.
7. The User Home Organization must be able to communicate information relevant to the session authorization via the Broker and the Service Provider to the User without allowing the Broker or the Service Provider to examine its content.

Nonrepudiation

8. There is a need for a recorded, authenticated and authorized agreement about the request for and delivery of service.

7. Computer Based Education and Distance Learning

This section describes the authorization aspects of computer based distance learning environments. In this section we will model the relationships and working practices in a hypothetical university environment where a student enrolls in courses, attends lectures, and takes the corresponding exams from remote locations (distance learning) or via computer equipment (computer based education). When completed successfully, a student is authorized to enroll in a set of subsequent courses according to his or her curriculum requirements. Completion of required courses with passing grades results in graduation.

Although this section specifically describes an example of a student taking courses at a faculty (department) of the university, the resulting requirements should also be valid for other applications in similar environments, e.g. library loans, electronic abstract and reprint services, computer and network access, use of copy machines, budget management, store retrievals, use of coffee machines and building access.

It is important to recognize that the AAA environment we are describing also needs to be managed. For example, for an application such as budget management, it is necessary to delegate budget authority from a central financial department to budget managers in education or faculty groups. An AAA environment must allow creation of policy rules either by certain individuals or by other AAA servers with authorization to do so.

7.1. Model Description

The establishment of the model involves four steps:

1. identification of the components that are involved and what they are called in this specific environment,
2. identification of the contractual relationships between the involved parties,
3. identification of the relationships that are based on trust, and
4. consideration of the sequence of messages exchanged between components.

7.1.1. Identification of Components

We will consider the components of a distance learning environment in the context of the conceptual entities defined in [2].

- The Student (User) -- the person enrolling in a course (Service) and taking the corresponding exam.
- The Educator (Service Equipment) -- the education content server for which the content is delivered by the Professor.
- The Educator Authorization Module (Service Provider AAA Server). This module must check at the service access point whether the student complies with the requirements for enrolling in the course. The authorization may be based on both local (by the professor) and remote policies (originating from the faculty). Rules must allow enough flexibility to prevent students from being falsely denied access to courses. Strict rules must only be applied at graduation time.
- The Faculty (Service Provider) -- the organization (department in U.S. terms) which controls the Service "Equipment" of which the Educator is one example.
- The Curriculum Commission (Part of User Home Organization) -- body responsible for creating rules by which a student is allowed to enroll in a certain course and how this course will count toward his or her graduation requirements. Students may legally take any course available at any time, however the Curriculum Commission will decide whether this course will contribute towards their graduation. When a Student registers with a certain Educator, the Educator may check with the Curriculum Commission AAA server whether the course will count towards graduation and confirm this with the student.
- The Student Administration (Part of User Home Organization) -- the administrative organization that authorizes students to enroll in courses if certain criteria, including financial criteria, are met. Next to the student, the Student Administration will keep track of any exam results for the student and will issue a graduation certificate when all criteria are met.

7.1.2. Identification of Contractual Relationships

Contractual relationships are illustrated in figure 19, below. Based on contract relationships, specific trust relationships are created as required.

Although not shown in figure 19, it is assumed that the university has contractual relationships with the faculties in which every faculty is allowed and obligated to build, maintain and present one or more specific studies.

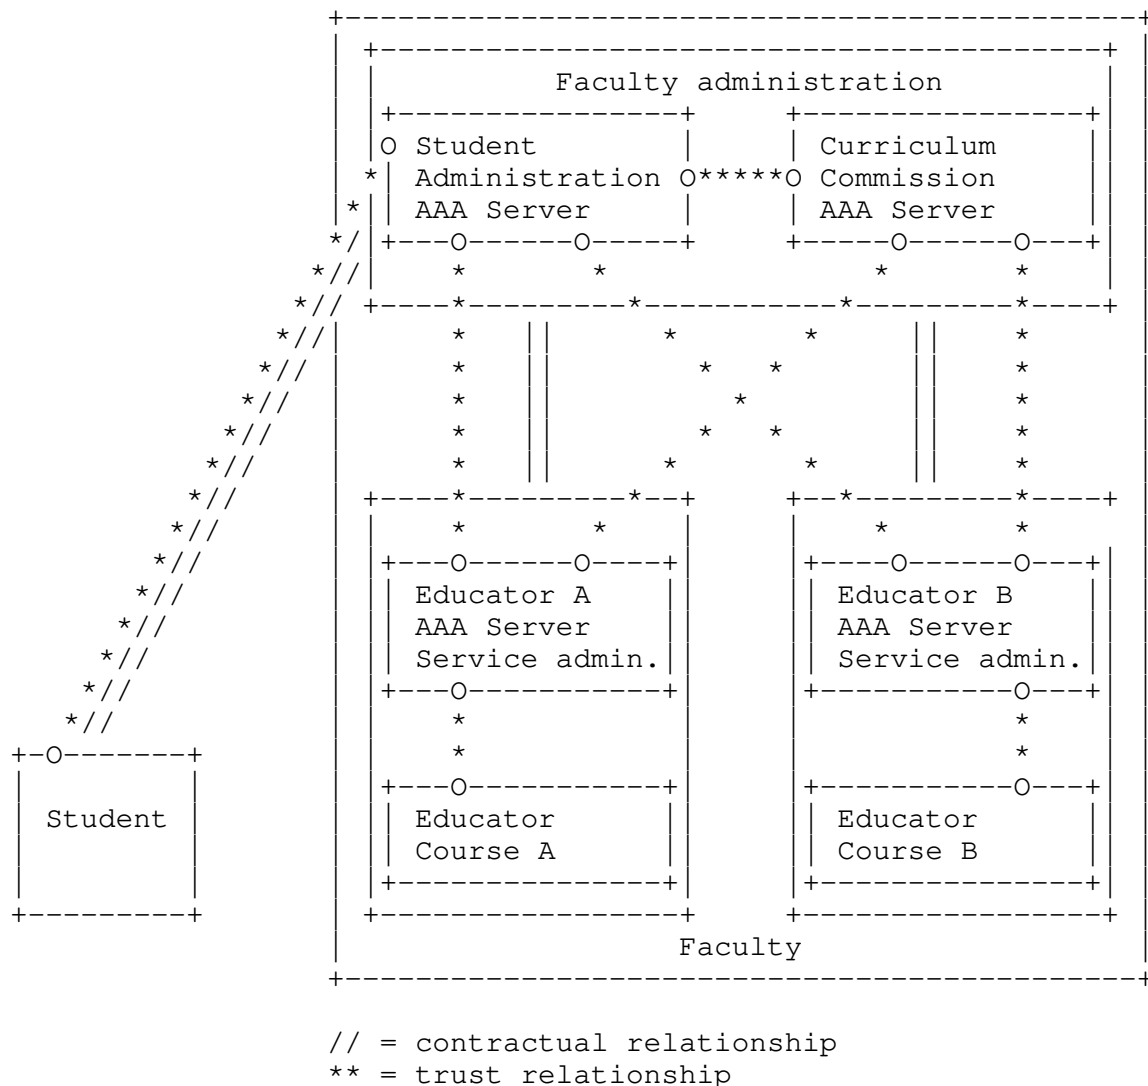


Fig. 19 -- Contractual relationships - single domain case

As shown in figure 19, the Student has a contractual relationship with the Faculty. The contract allows the Student to pursue a course of study consisting of a set of courses. Courses are presented to the Students by the Educators. A course of study may consist of courses from different Faculties.

Faculties have contracts among them allowing Students from one Faculty to enroll in courses from other Faculties.

Faculties instantiate Educators based on a contract between the Faculty Administration and the professor implementing and managing the Educator. Authorization is based on policy rules defined by one or more parties in the contractual relationships. For example, a professor has a policy to give the course only in the afternoon and the Faculty has a policy to give the course to their own students and students from faculty-x but not, when oversubscribed, to faculty-y students.

7.1.3. Identification of Trust Relationships

Figure 19 illustrates relevant trust relationships which statically enable AAA entities to communicate certain attributes in our simplified example. However, in order for the illustrated entities to work, other trust relationships that are not illustrated must already be in existence:

- A trust relationship based on a contract between the Faculty and the university enables a faculty to create and teach specific courses belonging to a course of study.
- Although not further detailed in this example, it is worth noting that trust relationships between faculties authorize students from one faculty to enroll in courses with other faculties.
- A professor responsible for the content of the Educator has a trust relationship with the administration of the faculty. Through this relationship, the faculty enables the professor to teach one or more courses fitting the requirements of the Curriculum Commission.

Figure 19 illustrates the following trust relationships:

- When a person wants to become a Student of a Faculty, the contract requires the Student to register with the Student Administration of the Faculty. If the requirements for registration are met, a trust relationship with the Faculty enables the Student to register for courses. For this purpose, the Student Administration will issue a student card which contains a student ID and information about the Faculty he or she is admitted to. The Student Administration will only admit Students who pay the necessary fees and have met certain prerequisites. The Student Administration will also keep track of Student grades and will ultimately issue a certificate at graduation. The Student Administration AAA server has access to relevant student data and will only issue grade information and other student-related information to authorized parties which have a specified means of authenticating.

- The Curriculum Commission AAA server needs a trust relationship with the Student Administration AAA server in order to obtain grade information to check whether a student has met the required course prerequisites. The Curriculum Commission creates certain rules within its AAA server which are evaluated when a particular student attempts to register for a particular course in order to give an advisory to the student.
- The Educator AAA server needs a trust relationship with the Student Administrator AAA server in order to verify whether this particular Student is in good standing with the Faculty. Only authorized Educator AAA servers may send requests to the Student Administration AAA server.
- The Educator AAA server needs a trust relationship with the Curriculum Commission AAA server in order to allow the Educator to obtain an advisory for the Student whether this course is consistent with his or her curriculum or whether the student meets the course prerequisites. Only authorized Educator AAA servers may send requests to the Curriculum AAA Server.

7.1.4. Sequence of Requests

For the sake of simplicity, we take the example of a student from the same faculty as the professor.

In this example the following interactions take place for a hypothetical course (see figure 20).

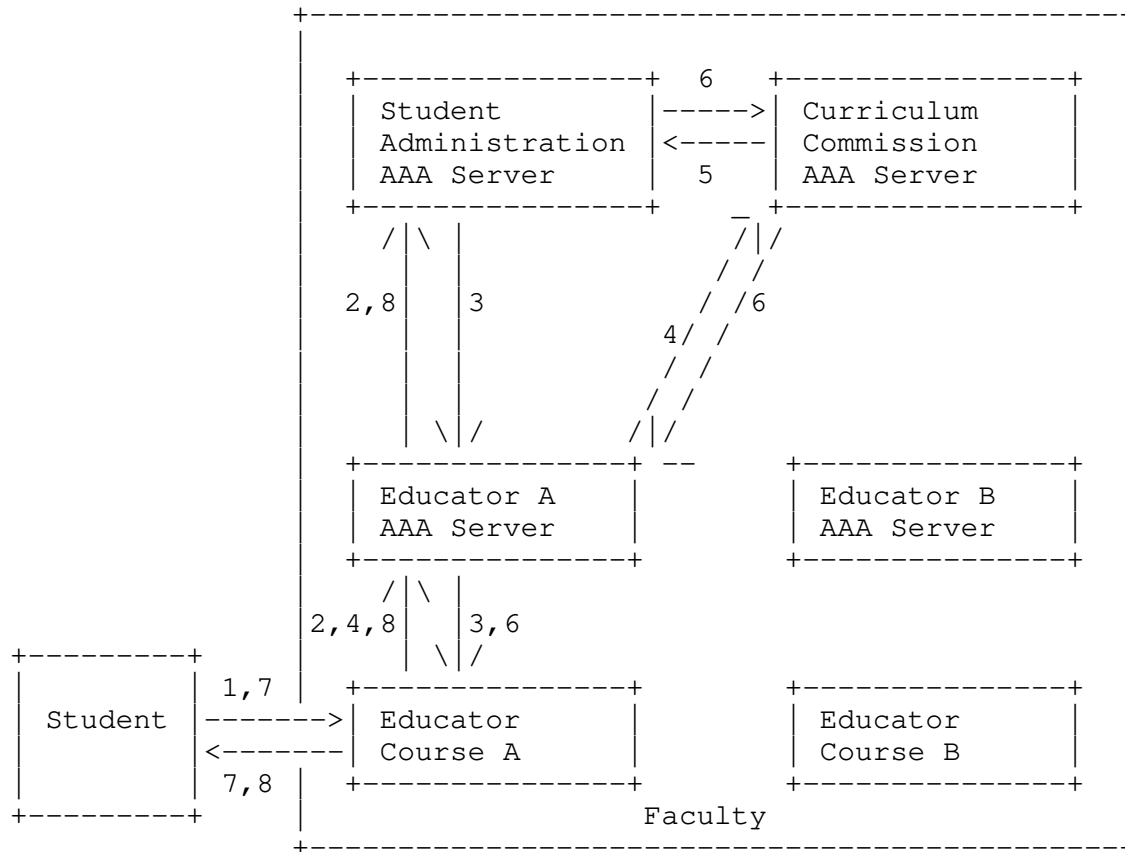


Fig. 20 -- AAA transactions - single domain case

1. After the Professor has set up the Service Equipment (Educator) students come to it presenting their ID (college card, name+faculty) and ask to be admitted to the course.
2. The Educator checks the ID to determine it is indeed dealing with a student from the faculty. This can include a check with the Student Administration.
3. The Student Administration replies to the Educator AAA Server, and the Educator AAA Server replies to the Educator.
4. The Educator checks the request of the Student against its own policy (courses only in the afternoon) and checks with the Curriculum Commission whether this student is advised to take the course. The necessary information is not normally known to or maintained by the professor.

5. The Curriculum Commission may check against the Student Administration to see if the Student had the necessary grades for the previous courses according to the policies set by the Curriculum Commission.
6. The Student Administration replies to the Curriculum Commission, the Curriculum Commission replies to the Educator AAA Server, and the Educator AAA Server replies to the Educator.
7. If now authorized, the Student is presented the material and the Student returns completed exams.
8. If the Student passes the tests, the Educator informs both the Student and the Student Administration that the Student has passed.

7.2. Requirements

We identify the following requirements for an AAA server environment for this example:

1. It must be possible to delegate authority to contracted partners. Although this requirement is not explicit in the limited example, the relationship between University and Faculty may require delegation of authority regarding the curriculum to the Faculty. In the case of budget management, this requirement is evident.
2. A system to manage the delegated authority must be established. It is possible that this is just another AAA server environment. This comes from the fact that one partner requires the presence of specific rules to be in the AAA server of another partner. For example, the Faculty must be sure that certain checks are performed by the Educator's AAA server.
3. AAA requests must either be evaluated at the AAA server queried or else parts of the request must be forwarded to another AAA server which can decide further on the request. As such, it must be possible to build a network of AAA servers in which each makes the decisions it is authorized to make by the relationships among the entities, e.g., a request from the Educator to the Curriculum Commission may result in a request to the Student Administration.
4. Transaction logs must be maintained to support non-repudiation for the grades of the students. This recording should be time-stamped and allow signing by authorized entities. A student should sign for taking an exam and this should be kept by the Educator's AAA

server. After grading, the professor should be able to sign a grade and send it to the Student Administrator and the Student Administrator's AAA server should log and timestamp this event.

5. Three types of AAA messages are required:

- authorization requests and responses for obtaining authorization,
- notification messages for accounting purposes, and
- information requests and responses for getting information regarding the correct construction of requests and for querying the database of notifications.

8. Security Considerations

The authorization applications discussed in this document are modeled on the framework presented in [2]. Security considerations relative to the authorization framework are discussed in [2].

Specific security aspects of each authorization application presented in this document are discussed in the relevant section, above.

Security aspects of the applications, themselves, are discussed in the references cited below.

Glossary

Attribute Certificate -- structure containing authorization attributes which is digitally signed using public key cryptography.

Contract Relationship -- a relation established between two or more business entities where terms and conditions determine the exchange of goods or services.

Distributed Service -- a service that is provided by more than one Service Provider acting in concert.

Dynamic Trust Relationship -- a secure relationship which is dynamically created between two entities who may never have had any prior relationship. This relationship can be created if the involved entities have a mutually trusted third party. Example: A merchant trusts a cardholder at the time of a payment transaction because they both are known by a credit card organization.

Policy Decision Point (PDP) -- The point where policy decisions are made.

Policy Enforcement Point (PEP) -- The point where the policy decisions are actually enforced.

Resource Manager -- the component of an AAA Server which tracks the state of sessions associated with the AAA Server or its associated Service Equipment and provides an anchor point from which a session can be controlled, monitored, and coordinated.

Roaming -- An authorization transaction in which the Service Provider and the User Home Organization are two different organizations. (Note that the dialin application is one for which roaming has been actively considered, but this definition encompasses other applications as well.)

Security Association -- a collection of security contexts, between a pair of nodes, which may be applied to protocol messages exchanged between them. Each context indicates an authentication algorithm and mode, a secret (a shared key, or appropriate public/private key pair), and a style of replay protection in use. [14]

Service Equipment -- the equipment which provides a service.

Service Provider -- an organization which provides a service.

Static Trust Relationship -- a pre-established secure relationship between two entities created by a trusted party. This relationship facilitates the exchange of AAA messages with a certain level of security and traceability. Example: A network operator (trusted party) who has access to the wiring closet creates a connection between a user's wall outlet and a particular network port. The user is thereafter trusted -- to a certain level -- to be connected to this particular network port.

User -- the entity seeking authorization to use a resource or a service.

User Home Organization (UHO) -- An organization with whom the User has a contractual relationship which can authenticate the User and may be able to authorize access to resources or services.

References

- [1] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.
- [2] Vollbrecht, J., Calhoun, P., Farrell, S., Gommans, L., Gross, G., de Bruijn, B., de Laat, C., Holdrege, M. and D. Spence, "AAA Authorization Framework", RFC 2904, August 2000.

- [3] Farrell, S., Vollbrecht, J., Calhoun, P., Gommans, L., Gross, G., de Bruijn, B., de Laat, C., Holdrege, M. and D. Spence, "AAA Authorization Requirements", RFC 2906, August 2000.
- [4] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [5] Aboba, B. and G. Zorn, "Criteria for Evaluating Roaming Protocols", RFC 2477, January 1999.
- [6] Beadles, Mark Anthony, and David Mitton, "Criteria for Evaluating Network Access Server Protocols", Work in Progress.
- [7] Aboba, B. and M. Beadles, "The Network Access Identifier", RFC 2486, January 1999.
- [8] Rigney, C., Rubens, A., Simpson, W. and S. Willens, "Remote Authentication Dial In User Service (RADIUS)", RFC 2138, April 1997.
- [9] Calhoun, P. and G. Zorn, "Roamops Authentication/Authorization Requirements", Work in Progress.
- [10] Perkins, C., "IP Mobility Support", RFC 2002, October 1996.
- [11] Glass, Steven, et al, "Mobile IP Authentication, Authorization, and Accounting Requirements", Work in Progress.
- [12] Hiller, Tom, et al., "cdma2000 Wireless Data Requirements for AAA", Work in Progress.
- [13] Neilson, Rob, Jeff Wheeler, Francis Reichmeyer, and Susan Hares, "A Discussion of Bandwidth Broker Requirements for Internet2 Qbone Deployment", ver. 0.7, August 1999, http://www.merit.edu/working.groups/i2-qbone-bb/doc/BB_Req7.pdf.
- [14] deBry, R., "Internet Printing Protocol/1.0: Model and Semantics", RFC 2566, April 1999.
- [15] Burdett, D., "Internet Open Trading Protocol - IOTP", RFC 2801, April 2000.
- [16] "SET Secure Electronic Transaction Specification Book 1: Business Description", Version 1.0, May 31, 1997, http://www.setco.org/download/set_bk1.pdf.

Authors' Addresses

John R. Vollbrecht
Interlink Networks, Inc.
775 Technology Drive, Suite 200
Ann Arbor, MI 48108
USA

Phone: +1 734 821 1205
Fax: +1 734 821 1235
EMail: jrv@interlinknetworks.com

Pat R. Calhoun
Network and Security Research Center, Sun Labs
Sun Microsystems, Inc.
15 Network Circle
Menlo Park, California, 94025
USA

Phone: +1 650 786 7733
Fax: +1 650 786 6445
EMail: pcalhoun@eng.sun.com

Stephen Farrell
Baltimore Technologies
61 Fitzwilliam Lane
Dublin 2
Ireland

Phone: +353 1 647 7406
Fax: +353 1 647 7499
EMail: stephen.farrell@baltimore.ie

Leon Gommans
Enterasys Networks EMEA
Kerkplein 24
2841 XM Moordrecht
The Netherlands

Phone: +31 182 379279
email: gommans@cabletron.com
or at University of Utrecht:
l.h.m.gommans@phys.uu.nl

George M. Gross
Lucent Technologies
184 Liberty Corner Road, m.s. LC2N-D13
Warren, NJ 07059
USA

Phone: +1 908 580 4589
Fax: +1 908-580-4991
EMail: gmgross@lucent.com

Betty de Bruijn
Interpay Nederland B.V.
Eendrachtlaan 315
3526 LB Utrecht
The Netherlands

Phone: +31 30 2835104
EMail: betty@euronet.nl

Cees T.A.M. de Laat
Physics and Astronomy dept.
Utrecht University
Pincetonplein 5,
3584CC Utrecht
Netherlands

Phone: +31 30 2534585
Phone: +31 30 2537555
EMail: delaate@phys.uu.nl

Matt Holdrege
ipVerse
223 Ximeno Ave.
Long Beach, CA 90803

EMail: matt@ipverse.com

David W. Spence
Interlink Networks, Inc.
775 Technology Drive, Suite 200
Ann Arbor, MI 48108
USA

Phone: +1 734 821 1203
Fax: +1 734 821 1235
EMail: dspence@interlinknetworks.com

Full Copyright Statement

Copyright (C) The Internet Society (2000). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

