Current Work on Authentication

Andrew D. Birrell, Butler W. Lampson,
Roger M. Needham and Michael D. Schroeder

Systems Research Center
Digital Equipment Corporation

We have been working on a design for an authentication service for a distributed system. The design has three goals that we feel have not been met simultaneously by any previous design. First, the service must be able to grow to cover an arbitrarily large physical area, arbitrarily many administrative organizations, and arbitrarily many users (millions or billions); the service must be suitable for a long lifetime. Second, the system must not be monolithically trusted: it must be possible to achieve authentication even if there exist untrusted parts of the system. Third, these goals must be met in such a way that each party to the authentication knows precisely what agencies the party must trust in order to accept the authentication.

The fundamental purpose of authentication is to enable *principals* to identify each other in a way that them to communicate, with confidence that the communication originates with one principal and is destined for the other. The principals we are considering include people, machines, organizations and network resources such as printers, databases or file systems.

Our design, like many others, is based on the use of a trusted intermediary known as a *key distribution center* or *authentication service.* In such a system it is necessary to be able to identify the principals. In a large scale system the only practical approach is for the principals to be named by some form of distributed name service.

A detailed description of our design for such a name service exists, although it is not yet in a form suitable for wide distribution (for pedagogic reasons, not completeness or correctness).

There are two interesting aspects to our design: its semantics, and our techniques for reasoning about the semantics.

Our semantics allow principals to achieve mutually assured authentication without necessarily requiring them to trust all the superior parts of the naming hierarchy.

Our authentication semantics are based on the notion of *role*, which is an authenticated path through the name space. Roles are suitable for inclusion in access control lists. Implicit with a role is the set of entities that must be trusted to authenticate the role. Use of different roles at different times allows a principal to choose what privilege he wishes to exercise.

Our reasoning starts with the observation that when two principals share a secret key, this forms a secure channel between them. This means that the two principals can readily use the key to establish communication encrypted by a conversation key, assured from freedom from eavesdropping, tampering or replay, and each assured of the identity of the other principal.

We view directories in the name service as being themselves principals. Each principal registered in the name service has a secure channel to that principal's directory (formed by the principals personal secret key or password).

Similarly, there are secure channels between pairs of directories, formed by shared secret keys.

We state what each party to a secure channel knows about the authentication process, and who he is trusting. We define a *forwarding rule* that permits formation of additional secure channels by composing existing ones. Our reasoning shows what each party to the newly formed channel knows and who he trusts. Our overall authentication algorithm consists of repeated application of the forwarding rule, to achieve a secure channel between two given principals by traversing a path between them across the name space.

Viewing the choice of this path as choice of *role*, produces the observation that we are authenticating roles and not identities. A principal indicates what role he wishes to exercise by nominating the path for the authentication algorithm.

Since our formal reasoning tells us about the initial knowledge and trust involved in a secure channel, and how these change when the forwarding rule is applied, we derive the knowledge and trust achieved by the complete authentication algorithm.

This method of composition of secure channels appears to be a powerful tool in designing authentication algorithms. Describing precisely the trust and knowledge at each stage allows us to have confidence in our algorithms (which has not been true of previous authentication algorithms, even from the same authors).