THIS PROVIDES A BRIEF OVERVIEW. WE WOULD LIKE TO WORK WITH YOU ON INCORPORATING ELEMENTS OF THE ATTRIBUTION FRAMEWORK OUTLINED HERE INTO OTHER GENI PROJECTS.

ATTRIBUTION FOR GENI

Jeffrey Hunker

Jeffrey Hunker Associates LLC
5109 Bayard St.
Pittsburgh, PA 15232
hunker@jeffreyhunker.com
202 257 7778

Matt Bishop

Dept. of Computer Science
University of California at Davis
1 Shields Ave.
Davis, CA 95616-8562
bishop@cs.ucdavis.edu
530 752 8060

Carrie Gates

CA Labs, Inc.
1 CA Plaza
Islandia, NY 11749-7000
carrie.gates@ca.com
631 935 2007

As people develop testbeds to facilitate the evolution of networks and network protocols, they can design support mechanisms for many forms of attribution, including those other than identity. Our project, "Attribution for GENI," has developed a set of requirements for attribution that will be useful both in the next generation infrastructure and in the data it manages.

Definition of Attribution: We define *attribution* as "the binding of data to an entity." So, for example, determining the identity of the sender of a message is attribution—binding data (the identity) to an entity (the sender). Similarly, attributing a delay in forwarding a packet to a particular network binds data (the length of the delay) to an entity (the particular network).

Some Motivation: Previous work on attribution in computer security rests on two basic assumptions: first, that the ability to attribute identity, or the property of interest, is beneficial; and second, that the needs of the various stakeholders are closely enough aligned that one can assume the needs of one (such as the security analysts) will satisfy all. Neither is in fact accurate.

As an example, Alice may wish her identity attributed when she connects to her bank's on-line web server to transfer money between accounts. In this case, both the sender (Alice) and the recipient (the bank) want to be able to attribute identity to one another. As another example, suppose that a government counterintelligence agent wants to access the web site of a terrorist organization. The web site may, or may not, want attribution. But the counterintelligence agent certainly will not want the terrorists to know that she is accessing (and possibly trying to break into) their web site. So, in this case, the sender (counterintelligence agent) does not want the recipient (terrorists) to be able to attribute anything.

Generalized Attribution Framework: This framework supports many different types of attribution, and recognizes that parties other than the sender and receiver may have an interest in the attribution choice. A negotiation system, or some other way of resolving the different requirements of different parties, would be desirable.

The different types of attribution are:

Perfect attribution, in which the binding of the data to the entity is known;
Perfect non-attribution, in which the binding of the data to the entity is unknown and undiscoverable;
Perfect selective attribution, in which the binding of the data to the entity is known to some set of entities, and unknown and undiscoverable by other entities;
Imperfect attribution, in which the binding of the data to the entity can be discovered, but doing so takes long enough that once discovered, the knowledge is useless or redundant, or discovering that

THIS PROVIDES A BRIEF OVERVIEW. WE WOULD LIKE TO WORK WITH YOU ON INCORPORATING ELEMENTS OF THE ATTRIBUTION FRAMEWORK OUTLINED HERE INTO OTHER GENI PROJECTS.

knowledge costs more than the value of knowing the attribution;
False attribution, in which the binding of the data to the entity appears to be known, but the attribution is incorrect but consistent over time;
Randomized false attribution, which is false attribution without the consistency over time; and
Unconcern, in which an entity does not care about the binding of the data to the entity.

Actors include the senders, the receivers, ISPs, their organizations (or governments), and backbone providers.

Attribution policies: Each actor has policies describing attribution requirements. Some way of reconciling policy conflicts is necessary—ideally an efficient policy negotiation structure, such as one involving an automated negotiation system as mentioned above.

The general framework looks like this:

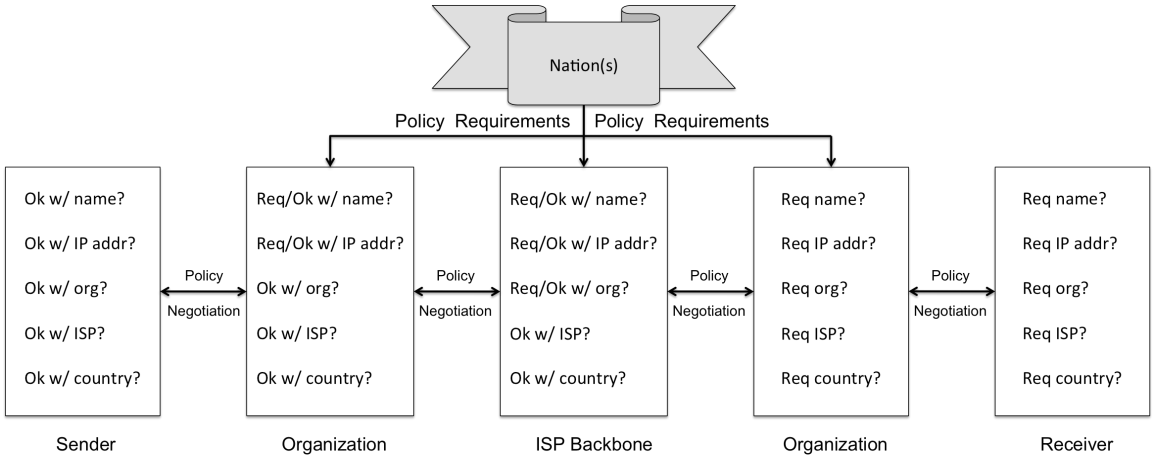

The framework focuses on five aspects of attribution:

The set of actors, which are chosen from the entities above;
The data (characteristic or property) being attributed, represented by a vector of values corresponding to the characteristics being attributed;
The assurance of the attribution, which is the confidence that the values being attributed are correct;
The entities to which the attribution is being provided, which may---or may not---be limited to the sender and recipient; and
The policy negotiation subsystem, which the entities use to negotiate the characteristics to be attributed and the assurance required for each, or to determine that the desired attribution or level of assurance is unobtainable.

Our full reports lay out the requirements for an attribution system. These are available at

http://nob.cs.ucdavis.edu/attrib